

Autonomic Link Management in Wireless Backhaul Networks with OpenFlow and Traffic Prediction

Maria Libunao, Bryan Ng and Winston K.G. Seah

School of Engineering and Computer Science, Victoria University of Wellington

P.O. Box 600, Wellington 6140, New Zealand

Email: marialibunao@gmail.com, {bryan.ng, winston.seah, liang.yang}@ecs.vuw.ac.nz

Abstract—Microwave radio is the predominant technology for backhaul or backbone networks, and its usage is expected to increase further. A way to increase capacity in this type of network is to aggregate multiple physical links into a single logical link. However, power is being wasted when the aggregated capacity is not being fully utilised. There are cost savings to be gained by only having the required number of physical links on. This paper presents a system that utilises extensions of the OpenFlow protocol to dynamically switch on/off physical links to meet capacity requirements while minimising power consumption while the Autoregressive Integrated Moving Average (ARIMA) model is used to predict the traffic load to be carried on aggregated (microwave) links. The proposed system is implemented on commercial-off-the-shelf microwave routers and experimentally validated on a testbed.

I. INTRODUCTION

The backhaul network links the access and core networks together. This can be realised through copper, fibre, and radio. Microwave radio is the predominant technology, making up 50% of backhauls in the world [1]. In fact, it is predicted that microwave usage will increase to 65% of cellular sites in 2020 [2]. Microwave radio has several benefits over the other two mediums, the most significant being cost and time for deployment. Both copper and fibre require trenches to be first dug up before the medium can be installed [3], and in areas with low consumer density the demand may be insufficient to justify the installation costs.

Networks in general, not only backhaul networks, are made up of many components such as routers and transmitters which all need electricity to operate. The amount of power consumed by Information and Communications Technologies (ICT) take up around 2-10% of the world's energy consumption [4]. This is expected to double by 2020. With a growing consciousness towards the effect of technology on the environment, there is no doubt that power consumption should be minimised.

Each microwave link is able to provide a certain amount capacity. In order to increase capacity, several links are used between routers. The links are then aggregated together into a single logical link [5]. However, the capacity provided by the aggregated links is not needed all the time, especially when traffic levels are low. In a backbone network, the link utilisation is typically around 30-40% [6]. Even in this scenario, the links are left on despite being idle and consuming energy.

Energy can be significantly saved when as few links as possible are kept on while still providing sufficient bandwidth to meet the traffic requirements. When the traffic load on a link increases, another physical link can be turned on. However, turning on a physical link is not instantaneous and can take between 8s and 20s before the link is operational, based on measurements done using commercial-off-the-shelf (COTS) devices. This implies the need to be aware of the impending increase in traffic loads early enough so that data packets are not lost during the powering up period of the new link.

The software defined networking (SDN) architecture provides logically centralised control of a network, with the ability to dynamically change the forwarding behaviour of network devices. SDN moves the control logic from the network devices to a logically centralised device, called the Controller, turning the network devices into simple forwarders that are programmable through a standardised protocol like OpenFlow [7]. This provides us with a viable platform to realise an autonomic link management system that dynamically turns on/off physical links on a microwave router to minimise energy consumption, yet provide sufficient bandwidth to meet traffic requirements, taking into account that a microwave link needs time to power up before it becomes operational.

In Section II, we discuss the background and related work on power saving techniques as well as traffic prediction algorithms. This is followed by the design of our system in Section III. After which, we introduce the ARIMA model in Section IV, followed by the experimental validation in Section V, before concluding this paper.

II. BACKGROUND & RELATED WORK

In this section, we review some proposed approaches for reducing energy consumption in backbone networks and methods/models used to estimate/predict the required link capacity, so that a decision can be made on the number of links to keep on or turn off.

A. Power Saving Techniques

The main techniques for power saving has been to either put components into a sleep state, or power them off. It has been shown that putting various components, e.g. line cards, crossbars, and main processors, to sleep can save power [8]. The decision can either be uncoordinated or coordinated. Uncoordinated sleeping is done per router using interarrival

times while coordinated sleeping is when routers collaborate together to determine sleeping time and aggregate traffic to a few nodes. Heller *et al.* [9] propose a technique to save energy in data centre networks. The decision is made through the traffic observed in the data centre network, by constantly evaluating the minimum number of components needed to support the traffic. To find the minimum subset, the optimal flow of the traffic must also be found. However, this is an NP-complete problem and unsuitable for realtime deployment. Similarly, Zhang *et al.* [10] discuss a method of re-routing for power savings for a wired network that uses OSPF [11] and MPLS [12]. Represented as a Mixed Integer Programming problem, the outcome depends on the set of nodes, links, and power consumption. A heuristic is used to reduce computation time where only a subset of paths, called the candidate paths, are used. As this solution only focuses on sleeping ports and line cards, it could only achieve a power saving on the line cards of 27% to 42%.

B. Capacity Estimation & Traffic Prediction

When links are bundled/aggregated between nodes, estimating the required capacity plays a critical role in deciding the number of physical links actually needed and how many can be put into sleep mode or even turned off. Capacity estimation often involves predicting the volume of traffic to be carried. Prediction models can be split into three categories: naive, parametric, and nonparametric [13].

Naive models are simple but not very accurate. Examples are instantaneous models which assume that traffic will be the same as the current traffic, averaging based on historical data and clustering similar patterns. Fukuda *et al.* [14] present an approach that samples the traffic load, then chooses the peak throughput during a set period. The empirical model is applied over a longer interval comprising multiple periods and samples to estimate the peak throughput of the next interval. Another approach, by Imaizumi *et al.* [15], uses a simple moving average to estimate the traffic load. Two nodes which are connected by a set of bundled links must agree on which links should be put into sleep mode as disabling one end would affect the capacity available in the entire aggregated link.

Parametric models, as the name implies, only need the parameters to be derived from the data, and the model structure is determined beforehand. Two widely used parametric models are traffic simulation models and time series models. Among time series models which predict future traffic based on past observations and an error term, the Autoregressive Integrated Moving Average (ARIMA) [16] model is applied to a wide range of areas, such as, road network traffic [17] and seasonal network traffic variations in a mobile cellular network [18].

Lastly, non-parametric models are more flexible with regards to parameters. However, the data must provide both the model structure and parameters. k -Nearest Neighbour, fuzzy logic, and neural networks are among the more common non-parametric approaches. Locally weighted regression is used by weighting each data point's residual and considering those which is the closest to the current state.

C. Wireless Transport using SDN

Traditionally, microwave backhaul networks have been managed by vendor-specific network management systems. The emergence of SDN has facilitated multi-vendor interoperability, and extensions to the OpenFlow protocol can lead to Software Defined Wireless Transport Networks [19]. The Wireless Transport project of the Open Networking Foundation (ONF) conducted three Proof-of-Concept (PoC) trials to demonstrate an SDN-based approach for managing a wireless transport network to break away from vendor proprietary NMS. The first PoC involved two concrete use cases: capacity-driven air interface and flow-based shaping [20], while the second [21] and third [22] PoCs demonstrated the capabilities and benefits of utilising a common Information Model for multi-vendor control of wireless network elements through open management interfaces. All these PoCs focused on multi-vendor interoperability and involved equipment from multiple vendors. The focus was not on the approach to dynamically manage the links based on traffic loads nor the model to estimate capacity or traffic prediction.

III. SYSTEM DESIGN

Muting a link [20] or putting it in sleep mode does not provide significant energy savings, as the circuitry remains on and continues to consume energy. The only way to reduce power consumption is to shut down a link completely, including all its associated circuitry. However, to turn it on again incurs a startup delay before the link can start carrying traffic. In this paper, we only focus on a group of aggregated/bundled links between two wireless (microwave) routers. When a link is turned off, the corresponding radios and associated circuitry at both ends (i.e. both routers) will be turned off to conserve energy. There are two key design components in our system: OpenFlow-based Link Management and traffic prediction, which we have chosen to use the ARIMA model in this case study.

A. OpenFlow-based Link Management

An agent process on the controller has two concurrent states: (i) polling routers/switches for traffic statistics and (ii) processing traffic statistics and predicting traffic loads, which include sending OpenFlow messages to network devices to turn the links on/off. The controller periodically polls the wireless routers/switches for the traffic statistics using a native OpenFlow message, viz. the Port Status message.

However, there were no existing OpenFlow message types to turn the links on/off or to report the amount of capacity available on the wireless routers. We use OpenFlow version 1.3 which has 226 unused message codes, we defined four new messages using these unused codes to accomplish the following tasks: (i) turn link on, (ii) turn link off, (iii) link change update, and (iv) capacity update.

When the data plane agent process receives OpenFlow messages, e.g. a Port Status message for "traffic statistics request", it will send traffic statistics back to the controller. Otherwise, it processes a packet as per OpenFlow semantics.

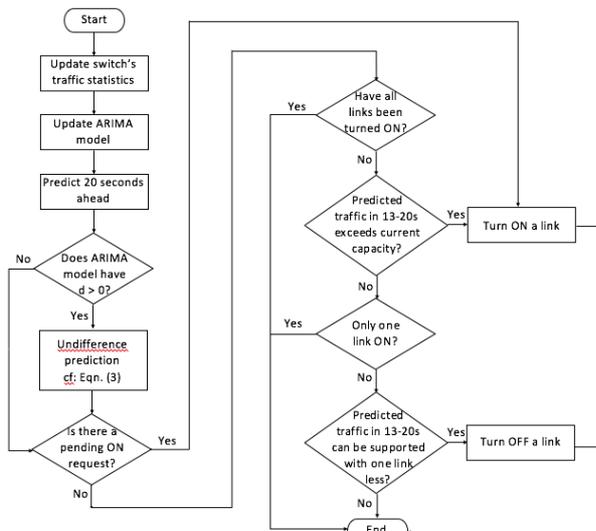


Fig. 1. A flow diagram of the prediction

B. Traffic Prediction

The flow diagram in Fig. 1 shows how the prediction is achieved and used to dynamically turn links on and off. The prediction is triggered upon receipt of traffic statistics from the switch. For our COTS device, a link can take as long as 20s to become fully operational after turning on, therefore we perform traffic prediction 20s ahead so that there is adequate time for the newly powered up link to meet the rising demand. A link will only be turned on if the traffic level 13 to 20 seconds later is predicted to exceed the current capacity where 13 seconds has been found to be the most often measured time. If the traffic is predicted to exceed the capacity within 13 seconds it will be too late to turn on a link, hence the need to predict 20s ahead to cater for the worst case. A link will only be turned off if the predicted traffic level can be supported by one less link, and single link is always left on because the radio transmitters need to exchange management frames.

IV. ARIMA MODEL

In our proposed system, we use the ARIMA time series model [16] to predict network traffic as it has been used for similar network-related scenarios [18]. Other prediction models can also be easily deployed on our system to suit other scenarios. An ARIMA model consists of the Auto Regressive (AR), Integrated (I), and Moving Average (MA) parts. The “AR” part uses historical data to produce predictions via regression of the data on itself, the “I” part is the order of differencing, i.e. how many times the data will be differenced (explained below) while the “MA” part uses errors from past prediction to improve subsequent predictions.

The ARIMA model is denoted by: $ARIMA(p,d,q)$, where p , d and q are the order for each component. This can be expressed as a time series:

$$y'_t = c + \phi_1 y'_{t-1} + \dots + \phi_p y'_{t-p} + \theta_1 e_{t-1} + \dots + \theta_q e_{t-q} + e_t \quad (1)$$

where y'_t is the data differenced d times, ϕ_i $\{i = 1, \dots, p\}$ are the parameters associated with the AR part of the model, with p representing the number of parameters while θ_i $\{i = 1, \dots, q\}$ are the parameters associated with the MA part of the model, with q parameters, and e_t are the forecasting errors.

Finding an appropriate $ARIMA(p, d, q)$ model to represent an observed stationary time series involves choosing the appropriate values for p and q (order selection) and estimation of the mean, the coefficients ϕ_i , θ_i , and the forecasting errors e_t ; for the detailed steps, we refer the reader to [16, Ch. 5]. Firstly, ARIMA requires the data to be stationary, i.e. not vary with time. The mean, variance, and covariance of the data must not be dependent on time. A technique called “differencing” developed by Box and Jenkins [23] is adopted to ensure that the time series is stationary. This is done through calculating the difference between each adjacent value. Eqn. (2) represents the differenced data of order 1, i.e. $d = 1$:

$$y'_t = y_t - y_{t-1}. \quad (2)$$

The parameter d represents the number of times differencing is needed to achieve data stationarity, which can be validated through unit root tests.

To determine the p and q orders, we examine the Autocorrelation Function (ACF) and Partial ACF (PACF). The ACF relates the values at y_t and y_{t-k} for some lag k while PACF removes the relation between the previous lags with the current lag since the current correlation might only have been a result of a previous correlation. Another way to choose a model is to compare the Akaike Information Criterion (AIC) [24] or the Bayesian Information Criterion (BIC) [25]. Both AIC and BIC use the likelihood that the model will produce the observed values. The model with the smallest AIC and BIC should be picked as the model.

A forecast can be obtained by inserting values into Eqn. (1). If the data had a non-zero differencing order, the forecast would be the lagged values. This means that only the difference between each value is predicted, instead of the actual value. A prediction without the differencing order can be obtained using a cumulative sum. The first value (y_1) is obtained from the actual observed data, and subsequent y_t values can be obtained using Eqn. (3):

$$y_t = y_1 + \sum_{i=1}^t y'_i \quad (3)$$

V. SYSTEM VALIDATION

We validated our system using a proof-of-concept prototype implemented using COTS hardware and ran experiments on a testbed, as shown in Fig. 2, to demonstrate the efficacy of the system in real world environments. Ryu is used as the SDN controller through which the control plane agent interacts with the other network devices and data plane agent. Two Aviat Networks CTR8540 microwave routers [26], each installed with four Outdoor Unit (ODU) radio transmitters, are connected as shown. The modulation scheme of the ODUs is configured to 1024QAM, giving each link an effective

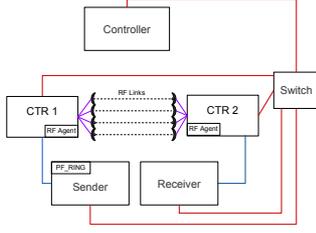


Fig. 2. Testbed Configuration

		Drop	Rise	Linear increase & decrease	Exponential increase & decrease
Packet loss	Prediction	1.98%	3.73%	0.26%	20.19%
	Threshold	19.26%	8.34%	0.02%	20.23%
Residual bandwidth	Prediction	32%	66.64%	35.16%	51.38%
	Threshold	28.26%	62.46%	47.69%	52.76%
Power consumption (kilowatts per hour)	Prediction	0.01996	0.02776	0.01674	0.01410
	Threshold	0.01544	0.02198	0.01993	0.01518

Fig. 3. Performance Comparison between Prediction and Threshold Methods.

throughput of 424Mbps. The Sender generates traffic using PF_RING™ [27] and sends it through a direct connection to microwave router CTR 1. The traffic is sent through the ODUs, across the microwave links, to the other microwave router, CTR 2, and out to the Receiver. On the Sender and Receiver, we run `tcpdump` [28] to collect packet level statistics.

A. Benchmark & Performance Metrics

We compare the performance of our approach against the threshold method, which was adopted in the capacity-driven air interface use case of the first ONF PoC [20]; in that study, the unused link was only muted but our approach turns the link off completely to save power. When the traffic levels exceed a certain threshold, a link will be turned on. Likewise if the traffic levels go below a certain threshold, a link is turned off. This particular method is static with no prior knowledge of historical data. Therefore, the thresholds remain constant even if the traffic rate may be ramping up at a different speed.

Since turning on a link that has been switched off incurs a startup delay, of between 8s and 20s, a hysteresis threshold method is adopted. The thresholds are set based on the link capacity (C_{link}) and the current number of active links (N_{links}). The threshold for a link to be turned on is given by:

$$Threshold_{ON} = (N_{links} - 1) \times C_{link} + \alpha \times C_{link}, \quad (4)$$

where α is the hysteresis factor and is typically used for tuning the sensitivity of threshold algorithms. In this paper, α is set to 0.8 for all our comparisons. For example, if there are two links on, a third link will be turned on if the traffic exceeds 763.2Mbps. Likewise, the threshold for a link to be turned off is given by:

$$Threshold_{OFF} = (N_{links} - 2) \times C_{link} + (1 - \alpha) \times C_{link} \quad (5)$$

With two links, the second link will be turned off if the traffic drops below 84.8Mbps.

The two methods are compared based on the following three metrics: (i) *packet loss*, (ii) *power consumption*, and (iii) *residual bandwidth*. Fig. 3 summarizes the performance results comparing the two methods over the four traffic scenarios, detailed in the following subsection.

B. Traffic Scenarios

1) *Periodic Drop in traffic every 30 seconds*: Traffic is sent at a constant rate of 700Mbps, with drops to zero every 30 seconds that lasts for 5 seconds. Such a traffic model is typical

in commute scenarios, like train or subway station, during peak periods. While waiting for the train to arrive, people browse the web, chat on social media, watch video, etc., and produce a steady flow of traffic for a nearby base station. But once the train arrives, there is a drop in traffic as the passengers depart from the base station. However, a new batch of passengers arrive after a while to wait for the next train. The chosen ARIMA model had an order of (0,1,1) and represented in Eqn. (6) and Eqn. (7). This means that the model only uses previous forecasting errors for the prediction.

$$y'_t = y_t - y_{t-1} \quad (6)$$

$$y'_t = -0.00883e_{t-1} + e_t \quad (7)$$

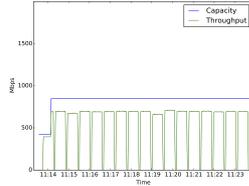
The prediction method keeps the links on as traffic is expected to resume after the drop, thus maintaining the throughput as shown in Fig. 4a. The threshold method, on the other hand, is quick to turn off a link whenever the traffic drops, and as a result packets are lost during the sudden rise since the capacity can no longer support the traffic.

From the power consumption aspect, the prediction approach consumes more power by not turning off the second link, as shown in Fig. 4b. Percentage of residual (unused) bandwidth is also higher (Fig. 4c) but this reduces packet loss by 17% compared to the threshold method.

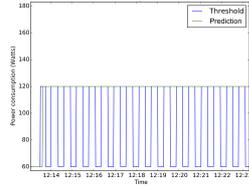
2) *Periodic Rise in Traffic*: In this scenario, traffic is sent over the RF links at a constant rate of 120Mbps for 10 seconds, with rises in traffic of up to 990Mbps, each lasting 5 seconds. Instead of people commuting home, consider a situation where people arrive for work. Once a train arrives at the station, there is a large surge of passengers alighting from the train that may cause a temporary rise in network traffic as the passengers alight. When the passengers leave the station, the traffic levels drop again until the arrival. The ARIMA model had an order of (4,0,4), as shown in Eqn. (8). There is no differencing order, so the model is based on the previous values and forecast errors.

$$y_t = 0.805y_{t-1} + 0.584y_{t-2} - 0.552y_{t-3} - 0.552y_{t-4} - 0.292e_{t-1} - 1.360e_{t-2} - 0.292e_{t-3} + 0.999e_{t-4} + e_t \quad (8)$$

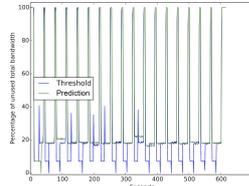
The throughput of the prediction model can be seen in Fig. 5a. Our system predicts that three links are needed for majority of the time and keeping them on to achieve lower packet loss. There are times where the number of links drop down to two. These are prediction errors that are fed back into the ARIMA model.



(a) Throughput & capacity

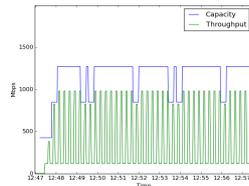


(b) Power consumption

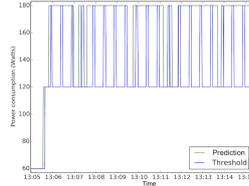


(c) Residual/unused bandwidth

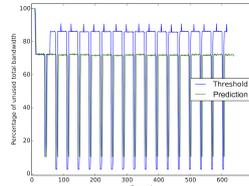
Fig. 4. Periodic Drop



(a) Throughput & capacity

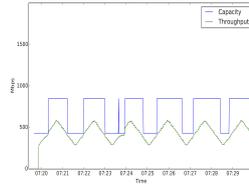


(b) Power consumption

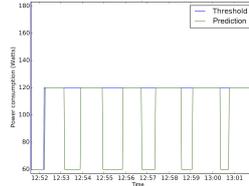


(c) Residual/unused bandwidth

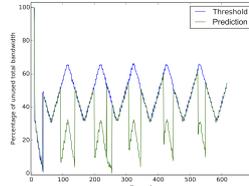
Fig. 5. Periodic Rise



(a) Throughput & capacity

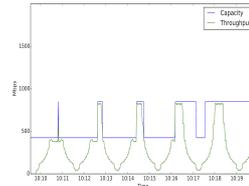


(b) Power consumption

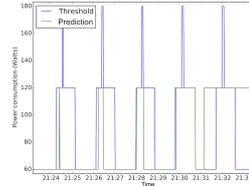


(c) Residual/unused bandwidth

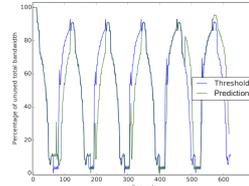
Fig. 6. Alternating Linear



(a) Throughput & capacity



(b) Power consumption



(c) Residual/unused bandwidth

Fig. 7. Alternating Exponential

On the other hand, the threshold method is quick to turn off the third link once the traffic drops. It then attempts to bring up another link once the traffic rises again. However, since the links take time to turn on, there is traffic loss for every other rise in traffic. Similar to the previous case, the threshold method has a higher packet loss, however it has lower power consumption (Fig. 5b) and residual bandwidth (Fig. 5c.)

3) *Alternating Linear Increase and Decrease:* This traffic model involves a linear increase of 2Mbps per second for 13 seconds, followed by a decrease of 2Mbps per second for a further 13 seconds, and repeating the cycle. The chosen ARIMA model had an order of (4,1,4), with an equation seen in Eqn. (9) where y'_t is the same as Eqn. (6).

$$y'_t = 0.146y'_{t-1} + 0.336y'_{t-2} + 0.006y'_{t-3} + 0.456y'_{t-4} - 0.528e_{t-1} - 0.649e_{t-2} + 0.072e_{t-3} + 0.477e_{t-4} + e_t \quad (9)$$

The received throughput can be seen in Fig. 6a. Our approach is able to predict the traffic going up and down, with a second link not being needed all the time. However, a prediction error occurs at around the 07:24 mark where the link is turned on too late and some packet loss occurs.

4) *Alternating Exponential Increase and Decrease:* The traffic model is similar to the previous case but involves an exponential increase in traffic level, followed by an exponential decrease. The chosen ARIMA model had an order of (0,1,4), denoted by Eqn. (10):

$$y'_t = -0.073e_{t-1} - 0.025e_{t-2} - 0.021e_{t-3} + 0.464e_{t-4} + e_t \quad (10)$$

Due to the fixed threshold, the second link remains on for the threshold method as the traffic level is not low enough to switch the link off. This results in higher power consumption (Fig. 6b) and more residual bandwidth (Fig. 6c.) The traffic patterns that show linear variation features favour the threshold method in terms of minimising packet loss. The lack of randomness does not benefit from the use of ARIMA.

Both methods give similar results for the three performance metrics, as shown in Fig. 7a to Fig. 7c. Interestingly, the behaviour of the prediction method gradually increases the amount of time the second link is turned on. This is based on the previous predictions which predicted an increase in traffic too late. Consequently, the packet loss decreased every time the traffic ramped up. This differs from the threshold method which leaves the links on for the same period of time, so the packet loss does not decrease as time goes on.

C. Results Analysis & Discussion

The proposed prediction method was compared against a threshold method, which uses a fixed threshold to decide when to turn links on and off. The prediction method achieved lower packet loss when traffic rate fluctuates rapidly, and provides lower power consumption and residual bandwidth in a gradually changing traffic conditions. Whenever the prediction method results in a lower percentage of packet loss, the difference between the prediction and threshold method is significant. However, a reduction of packet loss comes at the cost of having higher power consumption and residual bandwidth. Conversely, a reduction in power consumption

comes at the cost of having a higher packet loss. A balance of the two metrics was achieved through the proposed prediction method presented in this paper. In order to reduce packet loss, it attempts to provide enough capacity for future traffic levels by turning on the appropriate number of links in advance. To reduce power consumption, the prediction method turns off unnecessary links when it forecasts future traffic to be supportable by few than the currently active links.

VI. CONCLUSION

Microwave radio is the predominant technology used in wireless backhaul networks that link access networks to the core backbone networks. To increase the capacity between two wireless backhaul routers, multiple radio links are aggregated together to form a single logical link. However, power is often wasted when the traffic being carried across the network can be supported on a subset of links but all links are still kept on. These links should be turned off when not used.

The key contribution of our work lies in the use of the SDN approach to leverage the global view of the network, as well as, the logically centralised control. This allowed for a solution to be formulated using the OpenFlow protocol to monitor the network traffic through the global view, and use OpenFlow extensions to dynamically turn on and off links. The traffic levels are continually fed into an ARIMA model which predicted future traffic levels; other prediction models as well as machine learning techniques can also be used. As a newly turned on radio link (and its associated circuitry) takes time to power up and become operational, future traffic levels must be predicted early enough so that new radio links can be turned on and be ready to handle the increased traffic. Merely muting a link does not save much power as the circuitry continues to consume power.

This study has focused on a single aggregated/bundled link between two wireless routers. In order to minimise packet loss, a new link is turned on at the cost of increased power consumption and higher residual/unused bandwidth. From a network-wide perspective, instead of turning on a new link to carry the small volume of additional traffic, the SDN controller can re-route the additional traffic to links that have residual bandwidth, albeit very plausibly over a longer route/path that may incur a higher end-to-end delay for the re-routed traffic. This forms our ongoing and future work.

ACKNOWLEDGMENT

The authors thank Aviat Networks (NZ) Ltd for providing the equipment and technical support to carry out this research.

REFERENCES

- [1] T. Naveh. (2009, October) Mobile Backhaul: Fiber vs. Microwave. [Online]. Available: http://www.digitalairwireless.com/files/Fiber-vs-Microwave-White-Paper_1333235596.pdf
- [2] Ericsson. (2015, September) Microwave towards 2020. [Online]. Available: <https://www.ericsson.com/res/docs/2015/microwave-2020-report.pdf>
- [3] X. Huang, Y. J. Guo, A. Zhang, and V. Dyadyuk, "A multi-gigabit microwave backhaul," *IEEE Communications Magazine*, vol. 50, no. 3, pp. 122–129, March 2012.
- [4] S. Şener, E. Sardoğan, S. Staub, C. Eyupoglu, and M. A. Aydin, "World Conference on Technology, Innovation and Entrepreneurship Energy Efficiency in Backbone Networks," *Procedia - Social and Behavioral Sciences*, vol. 195, pp. 1966 – 1970, 2015.
- [5] S. Little, "Is microwave backhaul up to the 4G task?" *IEEE Microwave Magazine*, vol. 10, no. 5, pp. 67–74, August 2009.
- [6] W. Fisher, M. Suchara, and J. Rexford, "Greening Backbone Networks: Reducing Energy Consumption by Shutting off Cables in Bundled Links," in *Proc of the First ACM SIGCOMM Workshop on Green Networking*, New Delhi, India, 30 Aug - 3 Sep 2010, pp. 29–34.
- [7] Open Networking Foundation. OpenFlow. [Online]. Available: <https://www.opennetworking.org/sdn-resources/openflow>
- [8] M. Gupta and S. Singh, "Greening of the internet," *Proc of the 2003 conference on Applications, technologies, architectures, and protocols for computer communications*, pp. 19–26, 2003.
- [9] B. Heller, S. Seetharaman, P. Mahadevan, Y. Yiakoumis, P. Sharma, S. Banerjee, and N. McKeown, "ElasticTree: Saving Energy in Data Center Networks," in *Proc of the 7th USENIX Conference on Networked Systems Design and Implementation (NSDI)*, San Jose, California, USA, 28-30 April 2010, pp. 1–16.
- [10] M. Zhang, C. Yi, B. Liu, and B. Zhang, "GreenTE: Power-aware Traffic Engineering," in *Proc of the 18th IEEE International Conference on Network Protocols (ICNP)*, Kyoto, Japan, 5-8 Oct 2010, pp. 21–30.
- [11] J. Moy, "OSPF Version 2," Internet Requests for Comments, RFC Editor, STD 54, April 1998. [Online]. Available: <http://www.rfc-editor.org/rfc/rfc2328.txt>
- [12] J. Guichard, F. L. Faucheur, and J.-P. Vasseur, *Definitive MPLS Network Designs*. Cisco Press, 2005.
- [13] C. Van Hinsbergen, J. Van Lint, and F. Sanders, "Short Term Traffic Prediction Models," in *Proc of the 14th World Congress on Intelligent Transport Systems*, Beijing, China, 9-13 Oct 2007.
- [14] Y. Fukuda, T. Ikenaga, H. Tamura, M. Uchida, K. Kawahara, and Y. Oie, "Performance Evaluation of Power Saving Scheme with Dynamic Transmission Capacity Control," in *GLOBECOM Workshops, 2009 IEEE*, Nov 2009, pp. 1–5.
- [15] H. Imaizumi, T. Nagata, G. Kunito, K. Yamazaki, and H. Morikawa, "Power Saving Mechanism Based on Simple Moving Average for 802.3ad Link Aggregation," in *Proc of the IEEE GLOBECOM Workshops*, Honolulu, HI, USA, 30 Nov - 4 Dec 2009, pp. 1–6.
- [16] P. J. Brockwell and R. A. Davis, *Introduction to Time Series and Forecasting*, 3rd ed. Springer, 2016.
- [17] M. M. Hamed, H. R. Al-Masaeid, and Z. M. B. Said, "Short-term prediction of traffic volume in urban arterials," *Journal of Transportation Engineering*, vol. 121, no. 3, pp. 249–254, 1995.
- [18] Y. Shu, M. Yu, J. Liu, and O. W. Yang, "Wireless traffic modeling and prediction using seasonal ARIMA models," in *Proc of the IEEE International Conference on Communications (ICC)*, Anchorage, Alaska, USA, 11-15 May 2003, pp. 1675–1679.
- [19] D. Bercovich, L. M. Contreras, Y. Haddad, A. Adam, and C. J. Bernardos, "Software-Defined Wireless Transport Networks for Flexible Mobile Backhaul in 5G Systems," *Mobile Networks and Applications*, vol. 20, no. 6, pp. 793–801, 2015.
- [20] Open Networking Foundation, "Wireless Transport SDN Proof of Concept," White Paper, 9 October 2015.
- [21] Open Networking Foundation, "Wireless Transport SDN Proof of Concept 2," Detailed Report, 6 June 2016.
- [22] Open Networking Foundation, "Third Wireless Transport SDN Proof of Concept," White Paper, December 2016.
- [23] G. E. P. Box, G. M. Jenkins, and G. C. Reinsel, *Time Series Analysis Forecasting and Control*, 4th ed. John Wiley & Sons, Inc., 2008.
- [24] H. Akaike, "A new look at the statistical model identification," *IEEE Trans on Automatic Control*, vol. 19, no. 6, pp. 716–723, Dec 1974.
- [25] G. Schwarz, "Estimating the Dimension of a Model," *Ann. Statist.*, vol. 6, no. 2, pp. 461–464, 03 1978.
- [26] Aviat Networks. CTR 8540 MICROWAVE ROUTER. <http://aviatnetworks.com/products/microwave-routers/ctr8500-microwave-routers/>
- [27] ntop, "PF_RING™ – High-speed packet capture, filtering and analysis." [Online]. Available: http://www.ntop.org/products/packet-capture/pf_ring/
- [28] "TCPDUMP & LIBPCAP" [Online]. Available: <http://www.tcpdump.org/>