Weather Disruption-Tolerant Self-Optimising Millimeter Mesh Networks

Abdul Jabbar, Bharatwajan Raman, Victor S. Frost, and James P.G. Sterbenz

Information and Telecommunication Technology Center,
The University of Kansas, Lawrence, KS, USA
{jabbar,bharat,frost,jpgs}@ittc.ku.edu
http://www.ittc.ku.edu/resilinets

Abstract. Millimeter-wave networks have the potential to supplement fiber in providing high-speed Internet access, as well as backhaul for emerging mobile 3G and 4G services. However, due to the high frequency of operation (70–90 GHz), such networks are highly susceptible to attenuation from rain. In this paper, we present several mechanisms to overcome the disruptive effects of rain storms on network connectivity and service reliability. A resilient mesh topology with cross-layering between the physical and network layer has the capability to self-optimise under the presence of unstable links. We present a novel domain-specific predictive routing algorithm P-WARP that uses real-time radar data to dynamically route traffic around link failures as well as a modified link-state algorithm XL-OSPF that uses cross-layering to achieve resilient routing. Simulations are conducted to evaluate the effectiveness of the proposed algorithms.

Keywords: Resilient, survivable, weather disruption tolerant, millimeter-wave, mesh network, predictive routing, self-optimising.

1 Introduction

With the demand for high-bandwidth wireless communication growing at a rapid pace, coupled with the steady increase in the number of both fixed and mobile users, high-capacity backhaul networking solutions are needed that can quickly deliver large amounts of data as close to the end users as possible. Fiber-optic links provide the necessary bandwidth and reliability for both the backhaul and access links. However, there are two specific cases in which fiber-optic networks may be difficult or cost ineffective to deploy. The first case involves metropolitan areas in which the current fiber deployment is not adequate for the demand, either because existing fiber in a given location is at capacity, or because a particular service provider does not own fiber infrastructure and would have overcome significant deployment barriers. Secondly, in rural areas the subscriber density is generally inadequate to support the infrastructure needed for fiber-based broadband Internet access. With the average cost of new fiber ranging between $50,000 to $300,000 per kilometer, lower-cost fixed wireless links provide an increasingly attractive alternative.

K.A. Hummel and J.P.G. Sterbenz (Eds.): IWSOS 2008, LNCS 5343, pp. 242–255, 2008.

Fixed wireless deployments (e.g. 802.16 links in WiMAX service) are rapidly becoming an integral part of cellular networks that are continually pushing higher speeds towards end users with current 3G services and plans for 3.5 and 4G services. This has led to the current demand of higher bandwidth at lower cost. For commercial vendors, backhaul links must be as reliable as SONET circuits with 50 ms restoration times.[1]

However, until recently fixed wireless has had limited deployment due to three fundamental limitations:

1. *Limited capacity:* Traditionally, the capacity of the wireless networks is an order of magnitude or more lower than fiber.
2. *Reliability:* Wireless links are inherently error prone and lossy.
3. *Shared medium:* The shared medium of wireless networks adds complexity due to contention.

Fixed wireless mesh networks [3,4,5] are emerging that attempt to overcome these limitations and are successful to varying degrees [6]. Notably, 802.16 and other line-of-sight solutions are attempting to emulate wired links. However, the capacity of these links are still significantly lower than that of fiber.

More recently, millimeter-wave transmission [7,8,9] is being explored with frequencies of 71–86 GHz. Millimeter-wave links do not suffer from the capacity and shared-medium limitations: commercial radios can deliver data rates up to 1 Gb/s with the potential to go up to 10 Gb/s [10] using highly directional transmissions resulting in very narrow (pencil) beams. As a result, there can be several links originating or destined at a single node without causing interference, thus allowing for higher degrees of connectivity through spatial reuse. However, millimeter-wave networks (MMWNs) do not address the issue of reliability in wireless networks. On the contrary, millimeter-wave transmission exacerbates the problem of link stability due to its susceptibility to precipitation [7,11,12,13]. Therefore, it is essential for the wireless mesh networks as a whole to be disruption tolerant in the presence of rain storms. When the link experiences partial degradations or complete failure, the network should self-optimise in order to provide high reliability and availability [7].

In this paper, we investigate the effect of weather-related disruptions on millimeter-wave links and the effect of error-prone and failed links on the service availability of the network. Furthermore, we discuss a network-layer approach to overcome link instability by routing around failures within a mesh topology. This requires carefully designed cross-layer mechanisms to leverage physical-layer information at the network layer. We discuss two solutions to based on the link-state algorithm that optimise the network either reactively or predictively. First, we present XL-OSPF, a cross-layered version of the well-known OSPF routing protocol. This reactive approach uses link-cost metrics derived from

[1] While one can debate the merits of such a tight requirement for many data applications, these backhaul links are intended to be used as T1 and SONET replacements for which the SONET APS (automatic protection switching) 50 ms bound [1,2] is considered a hard requirement by at least some large mobile service providers.

physical-layer information to maximise the performance of OSPF in the presence of degraded links. Secondly, we present a *predictive weather assisted routing protocol* (P-WARP), that utilises short-term weather forecasts to reroute *ahead* of an impending link failure. We compare the performance of both approaches in terms of efficiency and the ability to self-optimise under weather disruptions. While some of the findings in this paper regarding disruptive effects of weather are well known, an important contribution is the quantitative characterisation of these challenges and their potential solutions.

1.1 Scope

The effect of a weather disruption on millimeter-wave networks mainly depends on two factors: the rain rate and the geographic footprint of the storm. Both of these parameters vary from one geographic region to the other. For example, the analysis of weather data from southern Great Plains region of the US [14] shows that approximately 78% of all storms are smaller than 25 km in diameter and account for only 1.0% of the precipitation. Furthermore, only 1% of the storms are large (over 40 km diameter) and account for 85% of precipitation. The remaining 20% are medium sized storms (20–40 km diameter) accounting for 14% of precipitation. We draw two conclusions from this study: First, the majority of the storms are small enough for a metropolitan size mesh network (approx. 1000 km^2) to reroute traffic around the storm. Secondly, even moderate sized storms are likely to have small size heavy intensity regions since they do not account for a significant percentage of rainfall. Extensive measurements over a 12-month period in the Topeka – Lawrence – Kansas City corridor confirm these conclusions. This research is most applicable to geographic locations whose climate fits the above-mentioned storm patterns.

The rest of the paper is organised as follows: Section 2 presents the theory behind the effects of weather on millimeter-wave band. Disruption tolerance at the network layer is discussed in Section 3 with XL-OSPF and P-WARP presented in Sections 3.1 and 3.2 respectively. This is followed by performance analysis in Section 4. Finally, a summary of our research and future work is presented in Section 5.

2 Effect of Weather on Millimeter-Wave Links

The primary means of weather disruption on millimeter-wave links is rain fade, which leads to signal attenuation [7,15]. This degradation is due to the effects of scattering and water absorption caused by the rain droplets present in the transmission path. Moreover, rain droplet shape can also affect the amount of attenuation experienced by the transmission. Specifically, if the individual rain droplet shape is more oblate than spherical, then the transmission will suffer substantial attenuation if it is horizontally polarised relative to signals that are vertically polarised. As a result, signal polarisation, center frequency of the signal, and the rain rate are the three major factors that can impact the attenuation of a millimeter-wave transmission.

There are several models that relate millimeter-wave transmission attenuation to rain rates. Most notable among these models are the Crane models [16] and the ITU-R P.530 recommendation [17]. These models employ probabilistic and empirical methods to calculate the attenuation for a given probability distribution of the rain rate. The relationship between the signal attenuation and the rain rate is given by [16]:

$$A(R_p, D) = \alpha R_p^{\beta} \left[\frac{e^{u\beta d} - 1}{u\beta} - \frac{b^{\beta} e^{c\beta d}}{c\beta} + \frac{b^{\beta} e^{c\beta D}}{c\beta} \right], \quad d \leqslant D \leqslant 22.5 \text{ km} \tag{1}$$

$$A(R_p, D) = \alpha R_p^{\beta} \left[\frac{e^{u\beta d} - 1}{u\beta} \right], \; 0 < D \leqslant d \tag{2}$$

where A is the signal attenuation in dB, D is length of the path over which the rain rate is observed, R_p is the rain rate in mm/h, α and β are numerical constants, and u, b, c, and d are empirical constants defined as:

$$u = \frac{\ln(be^{cd})}{d}, \quad b = 2.3R_p^{-0.17}, \quad c = 0.026 - 0.03 \ln R_p, \quad d = 3.8 - 0.6 \ln R_p$$

Given a weather event with a rain rate of R_p mm/hr and a millimeter-wave link of length D, the received signal strength is decreased by A dB. This increases the effective bit error rate that is a function of received signal strength. In summary, the effect of a weather disturbance on the millimeter-wave network is quantified by the bit error rate and subsequently by the packet error rate on individual links. For example, measurements conducted using off-the-shelf millimeter-wave radios[2] show that rain rates higher than 5 mm/hr result in 100% packet error rate [18].

3 Routing in Millimeter-Wave Mesh Networks

MMWNs are connected in a arbitrary grid topology (Figure 1), in which several access nodes (e.g A, B, and D) are sending data towards a node (e.g C) that has external network access. For generality, we show that the nodes may be connected with alternative links (such as fiber or low frequency wireless for redundancy), even though the selection of appropriate alternatives is not the focus of this paper (see Section 5). Given this scenario, the objective of the MMWN routing protocols is to route around link failures without loosing data packets.

Earlier research has considered several different protocols for routing in mesh networks [19,20]. The focus of the community has been on issues related to shared medium contention and mobility, as well as metrics to quantify their effect on data transmission [21,22,23]. However, MMWNs are neither mobile nor do they have contention issues. Thus, with the exception of link instability, MMWNs share most of their characteristics with *wired* networks.

[2] 72.5 GHz radio with transmit power of 17 dBm, antenna gain of 43dBi, BFSK modulation, and Reed-Solomon FEC code over 10 km link using 1000B frames.

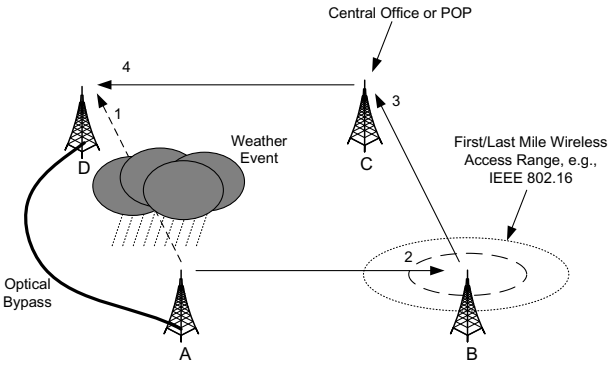

Fig. 1. Schematic of an MMWN with weather system

Conventional wired routing protocols, on the other hand, assume stable end-to-end paths composed of highly reliable links. Millimeter-wave links experience a continually varying link quality (state), particularly during rain storms, which leads to unstable end-to-end paths over time. Since the commonly used link-state protocol OSPF [24] does not mandate or recommend specific link metrics, the default implementations do not instrument explicit mechanisms in the protocol to support cost metrics based on physical link characteristics such as BER. Realising the need for OSPF to handle dynamic network scenarios, several modifications to support mobility are being currently reviewed [25,26]. However, in case of static MMWNs, the network dynamics are a result of link disruptions rather than mobility. In summary, the static, point-to-point, high-capacity but unreliable links of MMWNs present a unique case that can benefit from a domain specific network solution.

Due to the constantly varying link quality, it is necessary for the system to *self-optimise* in order to deliver reliable service. Thus, it is the goal of this paper to explore intelligent cross-layer mechanisms that optimise network routes so as to deliver the highest possible service availability and reliability.

3.1 XL-OSF: Cross-Layered OSPF

A cross-layer approach to link metrics could significantly improve the performance of dynamic link state algorithms [27,28] We chose OSPF (Open Shortest Path First) [24] because it has been widely researched and deployed. OSPF has two basic mechanisms to determine link state. One is the Link State Advertisements (LSAs) generated by each node carrying the status of all its links along with their costs, which are flooded throughout the network. Secondly, hello packets are used to determine if the link to a given neighbor is still alive. A *dead interval* based on the *hello interval* is used to detect dead links. The routes are reactively updated after the LSAs propagate through the network. With rapidly varying link quality, the only mechanism through which OSPF in its default configuration can detect link degradations is when four consecutive hello

packets are dropped. Since the size of the hello packets is much smaller than data packets, a BER that results in four consecutive hello drops will correspond to a significantly higher data packet drop rate.

One mechanism that can be used to improve the performance of OSPF is a cost metric that is proportional to the bit error rate of the link. However, this is a difficult proposition given the lack of information exchange between the link that detects the actual packet losses and the network layer that determines the routes. Several mechanisms are possible that use in band (packet header) or out-of-band (probe packets) signaling to determine the actual packet error rate.

For the purpose of analysis, we assume such a mechanism that informs the routing layer of the effective packet error rate (PER). We define a cost metric that is proportional the the effective PER, carefully choosing the scaling factor so as to avoid route fluctuations. The exact values of the scaling factor and link weights are given in Section 4.1. The performance of OSPF with this cost metric is discussed in Section 4.2.

Even with the error-based cost metrics, OSPF remains a reactive protocol that requires a finite amount of time before it adapts to changes in link state. If the application or service demands a highly-reliable service, the reactive protocols must have a very short update interval on the order of milliseconds. But this adds unacceptable level of overhead even for broadband networks as discussed in Section 4.2. In the following section, we discuss a predictive alternate for routing.

3.2 P-WARP: Predictive Weather Assisted Routing Protocol

As discussed above, reactive algorithms cannot meet stringent 50 ms restoration service requirements in MMWNs during weather disruptions. Furthermore, it is difficult to measure effective BER or PER at the routing layer without an explicit signaling mechanism. In this section, we investigate the use of information *external* to the network in order to predict the future state of links.

The proposed predictive routing algorithm is a link-state algorithm that utilises weather radar to forecast the future state of the link. In contrast to the OSPF solution mentioned previously, the primary difference is the mechanism through which the link costs are obtained. While XL-OSPF depends on BER measurement from errored packets, we propose P-WARP (predictive weather assisted routing protocol), in which BER of each link is calculated from weather data that is modeled in real-time using the following methodology:

Modeling Effective Rain Rate from Weather Radar. Short-term weather prediction obtained from doppler radar is used to model the predicted rain in the geographical region of the mesh network. Radar data is first passed through a geometric model that maps the areas corresponding to different rain intensities to ellipses. For the sake of simplicity, all precipitation regions are divided in to 3 categories representing heavy, moderate and light, or no rain. The rain rates assigned to each category depends on the characteristics of transceivers such as link budget and the length of the transmission link. For example, measurements on a typical off-the-shelf millimeter-wave system suggest that for a 10 km link,

heavy rain corresponds to an average rate of 5 mm/hr or more, moderate rain corresponds to 2–5 mm/hr, and light or no rain corresponds to an average rate less than 2 mm/hr.

Link Cost from Effective Rain Rate. The rain rate obtained in the previous step is then used to determine the effective BER on all the links of the network based on calculations from Section 2. This processing is done at a either a single *core node* or a small subset of nodes that are connected to the external network and are capable of receiving weather-radar data. The topology and physical location of the network is pre-programmed into the software module that does the link BER calculation as well as the PER (packet error rate) for a predefined average packet size. Finally, the link costs are calculated from the PER as discussed in Section 4.1.

Link Status Updates and Route Computation. The weather-based link status updates (WLSUs) in P-WARP are slightly different from the link state advertisements (LSAs) of OSPF. WLSUs are generated from the core nodes and contain the costs of all links in the network based on their *predicted* quality. These weather based updates are flooded throughout the network. When an individual node receive a WLSU, it recomputes routes using the Dijkstra shortest path first algorithm. However, unlike OSPF individual nodes do not generate separate LSAs for the links to their neighbors. This approach significantly reduces the protocol overhead because only *one* update is generated for *all* the links and updates are generated only when a change in one or more link costs is predicted. Thus, the network reroutes traffic *ahead* of the incoming storm thereby minimising the packet loss.

Route Sensitivity. Due to storm movement, the effective BER on each link will vary continuously over the duration of the event. In order to avoid route flaps and false alarms, we use thresholds along with hysteresis. A minimum noticeable change $\mathrm{BER}_{\mathrm{thresh}}$ is defined below which all BER changes are ignored, which determines the resolution of the algorithm. Further, a hysteresis percentage H_{thresh} determines the minimum change in the cost of a link for an update to be generated.

4 Performance Evaluation

In this section, we present our methodology for modeling and simulating storms. We also present the specific details of XL-OSPF and P-WARP which were used in the simulations. Finally, we quantify the disruptive effect of storms as well as the disruption tolerance of the proposed mechanisms.

4.1 Methodology

Storm Modeling. A typical example of a rain storm snapshot over part of the mesh network is shown in Figure 2(a) as a radar reflectivity map. The link length

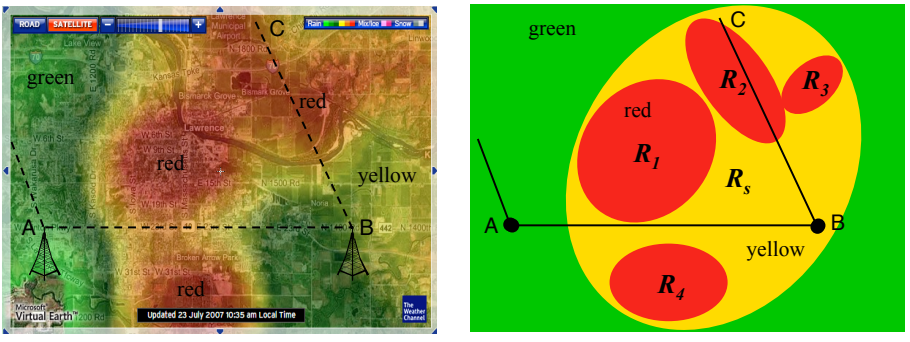

(a) Radar image (b) Geometric model

Fig. 2. Example storm model

in this example is 10 km. The geometric model of the storm at that instant is shown in Figure 2(b). Our model divides the radar reflectivity image in to regions of high (red), low (yellow) and little or no (green) precipitation. Each region is modeled as an ellipse (R_1–R_5) and individual links are represented as line segments. Equation 2 is used to calculate the effective BER and PER based on the intersection between the links and storm ellipses.

This example shows that while certain links (e.g. $\overline{\text{BC}}$) are severely degraded due to the heavy rain, it is possible to re-route traffic on other adjacent links (e.g. $\overline{\text{BA}}$) that are not in an intense rain region. Recall from Section 1.1 that this pattern of rain intensity distribution is typical for a majority of storms in the US Great Plains. In order to model a storm in real time, we continuously calculate the link attenuation and generate WLSUs as the storm moves through the network. The rate of WLSU generation is dependent on the dynamics of the storm relative to the network.

Routing Protocol Configuration. We evaluate the performance of XL-OSPF and P-WARP and compare them against OSFP using common defaults, as well as static routing as the lower bound on performance. For standard OSFP, we use the default values of 10 and 40 seconds for *hello* and *dead* intervals respectively [29]. The LSAs in XL-OSPF are rate limited to 10 seconds, whereas WLSUs in P-WARP are rate limited to 30 seconds, predicting the weather approximately a minute in advance.

For both XL-OSPF and P-WARP the link costs and other protocol parameters are set as follows. Assuming uniform distribution of the bit errors, the cost of a link between two nodes i and j is calculated as:

$$C_{ij} = P \times \text{BER}_{ij} \times \gamma \tag{3}$$

where P is the average packet size on the network, BER_{ij} is the bit error rate observed on the link, and γ is the scale factor. The scale factor determines the sensitivity of the link cost with respect to change in BER and is set to 1000 in

our simulations. A $\text{BER}_{\text{thresh}}$ of 10^{-8} is used to define the minimum observable change in BER with hysteresis using H_{thres} of 10% to avoid excessive route flaps in the network. Finally, the value of cost is bounded in the range of $[1, 1000]$ which determines the maximum number of hops a packet can traverse in order to avoid an error-prone or lossy link. Since, the objective here is to avoid disrupted links at all costs, we set this range to 1000.

4.2 Simulation and Results

We conducted simulations using storm modeling software that we developed in Matlab and the ns-2 simulator [30].

Simulation Setup. The simulated topology consists of 16 nodes connected in a square mesh as shown in Figure 3. The millimeter-wave links between any two nodes are 10 km long. The nodes remain fixed at their locations throughout the simulation. Modeling a cellular backhaul network, nodes 0 and 15 are connected to the external network (Internet) and hence are sink nodes. The remaining 14 nodes generate traffic at a rate of 2.4 Mb/s that is destined to one of the two sink nodes, picked randomly. The generated traffic is CBR (constant bit rate) over UDP with a packet size of 1000 bytes. We use two different storms to evaluate the disruption tolerance of the proposed mechanisms. These storms consist of an outer yellow ellipse varying between 20–30 km in diameter and inner red ellipses with a diameter varying between 5–10 km. The first storm consists of four high-intensity regions whereas the second storm contains two high-intensity regions. Both storms last for a duration of 30 minutes over the mesh network. We evaluate the packet delivery ratio and the service availability of the network for four different routing mechanisms.

Packet Delivery Ratio. The packet-delivery ratios averaged over a window of two seconds are shown in Figure 4 for all four routing protocols. This plot shows the instantaneous response of the network to the first simulated storm. As the individual links fail due the storm, the delivery ratio of the network falls rapidly. The time taken by the network to recover from link failures depends on

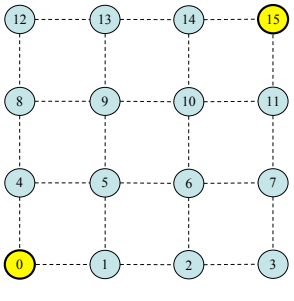

Fig. 3. Simulation topology

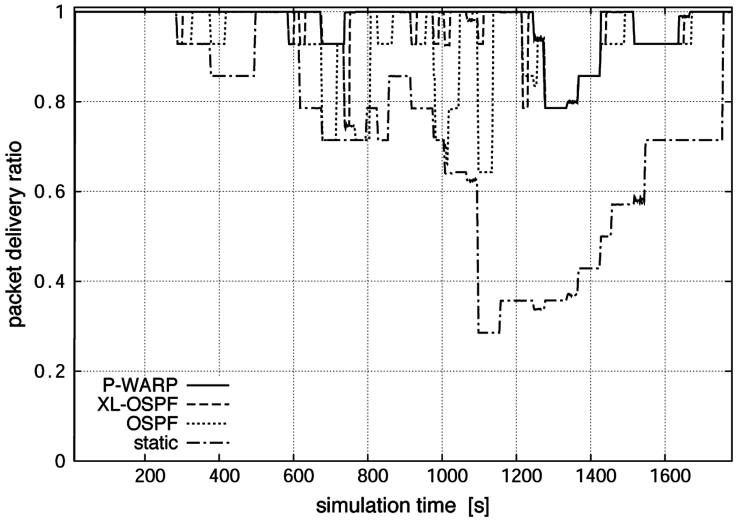

Fig. 4. Windowed average of received packets: first storm

the routing protocol. Static routing, as expected, performs very poorly marking the lower bound. OSPF without any modification performs better than static because it can sense link outages from the loss of four consecutive hello packets. However, the delay in detection and route re-computation results in significant packet loss. XL-OSPF performs better than static and standard OSPF because it can detect degrading links in a shorter time from the their cost, as advertised in the link state updates. Since the cost metric is directly proportional to the link error rate, high PER paths are avoided whenever possible. P-WARP outperforms all three protocols because it can predict an upcoming link failure from weather updates and reroutes traffic ahead of the disruption.

For example, consider the packet delivery ratios at $t = 1100$ s in Figure 4. At 1100 seconds, the storm disrupts several links causing severe packet loss in the case of static routing. Furthermore, the network does not recover until the storm passes at 1600 seconds. OSPF, on the other hand, detects failed links and recovers at 1140 seconds, indicating a 40 sec recovery time which corresponds to the dead interval. XL-OSPF recovers much faster compared to OSPF and static routing. It takes approximately 10 seconds to get back to 100% delivery ratio. Finally, P-WARP does not drop any packets at all confirming its predictive behavior. Accurately predicting the impending disruption, P-WARP preemptively routes data on stable paths, thereby avoiding the failed links completely.

In order to compare the aggregate performance of the protocols, the cumulative average of the packet delivery ratio is shown in Figure 5. The cumulative average of packets delivered by XL-OSPF is very close to that of P-WARP, both of which outperform conventional OSPF. In the case of XL-OSPF, the frequency of LSAs determines the reaction time of the network; strict restoration times would require very short update intervals. Because of its predictive nature, P-WARP has two distinct advantages: first, it has no reaction time which

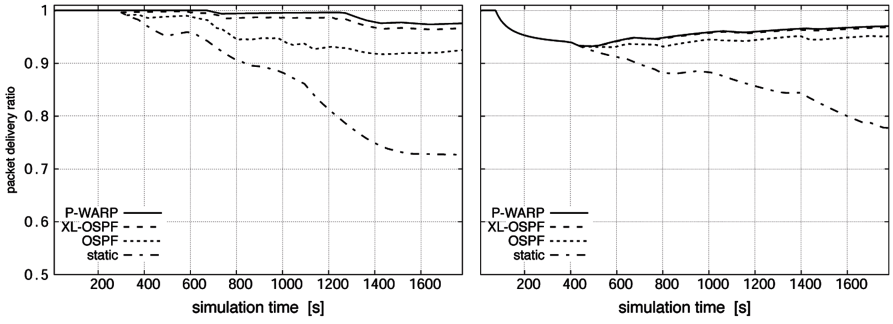

(a) First storm (b) Second storm

Fig. 5. Cumulative average of received packets

might be necessary if stringent service requirements such as 50 ms restoration are to be met; secondly, the frequency of weather updates does not scale with the restoration times, leading to lower protocol overhead.

Service Availability. In order to evaluate the availability of the network we consider a media streaming-application. The availability of the network is then defined as the percentage of time that the network can deliver data reliably to the destination with *no errors* within a given delay bound. Table 1 shows the availability of the network for the duration of the simulation for two different weather events. We selected a delay bound of 4 ms that is derived from cellular backhaul deployments. Predictive routing outperforms all the other cases since it routes traffic ahead of time, thus increasing the duration during which the service remains available. Among the other three mechanisms, XL-OSPF performs closest to predictive routing.

We have conducted several more simulations with varying update intervals for OSPF and XL-OSPF, and have observed that while the general relationship between the OSPF, XL-OSPF and P-WARP routing remains the same, the performance gap between XL-OSPF and P-WARP decreases with increasing LSA frequency, as expected.

Overhead. Static routing does not generate any overhead traffic. Conventional OSPF generates periodic hello messages as well as LSAs. However, the frequency

Table 1. Service availability

Protocol	Availability Storm-1	Availability Storm-2
P-WARP	0.9700	0.9638
XL-OSPF	0.9666	0.9554
OSPF	0.9514	0.9209
Static	0.7769	0.7304

of the LSA does not affect the performance because the quality of the link is not reflected in the its cost metric. On the other hand, XL-OSPF uses a cost metric that is proportional to link BER and hence generates frequent updates that are flooded in the network. In order to react quickly to weather disruptions, XL-OSPF must generate LSAs at a higher rate, leading to a significant increase in overhead. The number of updates generated in P-WARP is comparatively lower for two reasons. First, a single WLSU update carries the predicted costs for all the links. Hence, individual nodes do not generate link state updates. Secondly, an update is generated only when there is a change in the predicted BER of one or more links. Since the performance of the P-WARP is not dependent on the arbitrary update frequency, WLSUs are rate limited to one in every 30 seconds, representing the weather-prediction interval.

5 Conclusions and Future Work

In this paper, we have shown that millimeter-wave mesh networks must be self-optimising in order to provide reliable service under weather disruptions. It is observed that conventional reactive routing mechanisms do not perform well without a cost metric that reflects the physical status of the link. We presented two domain-specific mechanisms that utilise cross-layering between physical and network layer to improve the availability of network. In XL-OSPF, a cost metric that is proportional to the BER is shown to perform well under lightly loaded conditions. A novel predictive routing algorithm P-WARP based on a short-term weather forecast outperforms reactive routing both in terms of throughput and overheard. The proposed algorithm finds the optimal paths with relatively low frequency of routing updates. Simulations show that MMWNs with the proposed self-optimising mechanisms form a viable low cost solution to high speed backhaul and access networks.

We are currently evaluating the proposed schemes with real storm data collected by the US National Weather Service. Preliminary results based on real storms indicates a good match with the results shown in this paper. We intend to study several more cases to evaluate the performance of the network for different type of applications as well as different type of storms such as tropical, cyclonic. A fiber or low-frequency wireless bypass to overcome node failures remains a part of our future work. In order to support heavily loaded networks, we are investigating metrics that consider link quality as well as available capacity of the active links in its forwarding decisions. Future work also includes image processing of radar data to automatically calculate link attenuation on a linear scale.

Acknowledgments

We would like to acknowledge the meteorological guidance provided by Donna Tucker, the help with illustrations from Dan DePardo, and the technical support from Adam Hock. We also acknowledge the attenuation modeling work done by

Alex Wyglinski and Vinay Muralidharan. This project was funded in part by Sprint; we acknowledge the technical contributions and guidence of Tim Euler.

References

1. ANSI T1.105.01: Telecommunications – synchronous optical network (SONET) – automatic protection switching. American National Standard T1.105.01-2000 (March 2000)
2. ITU-T G.841: Types and characteristics of SDH network protection architectures. ITU-T Recommendation G.841
3. Uchida, D., Sugita, M., Toyoda, I., Atsugi, T.: Mesh-type broadband fixed wireless access system. NTT Technical Review 2(1), 44–54 (2004)
4. Wu, Y., Hui, J., Sun, H.: Fast restoring Gigabit wireless networks using a directional mesh architecture. Computer Communications 26, 1957–1964 (2003)
5. Whitehead, P.: Mesh networks: A new architecture for broadband wireless access systems. In: IEEE Radio and Wireless Conference, Denver, CO, USA, pp. 43–46 (September 2000)
6. Akyildiz, I.F., Wang, X., Wang, W.: Wireless mesh networks: A survey. Computer Networks 47(4), 445–487 (2005)
7. ITU-R F.1704: Characteristics of multipoint-to-multipoint fixed wireless systems with mesh network topology operating in frequency bands above about 17 GHz. ITU-R Recommendation F.1704 (2005)
8. Khan, J.A., Alnuweiri, H.M.: Traffic engineering with distributed dynamic channel allocation in BFWA mesh networks at millimeter wave band. In: Proceedings of the 14th IEEE Workshop on Local and Metropolitan Area Networks, Chania, Greece, pp. 1–6 (September 2005)
9. Ohata, K., Maruhashi, K., Ito, M., Nishiumi, T.: Millimeter-wave broadband transceivers. NEC Journal of Advanced Technology 2(3), 211–216 (2005)
10. Torkildson, E., Ananthasubramaniam, B., Madhow, U., Rodwell, M.: Millimeter-wave MIMO: Wireless Links at Optical Speeds. In: Proceedings of the 44th Allerton Conference on Communication, Control and Computing, Monticello, Illinois, USA (September 2006)
11. Izadpanah, H.: A millimeter-wave broadband wireless access technology demonstrator for the next-generation internet network reach extension. IEEE Communications Magazine, 140–145 (September 2001)
12. Hendrantoro, G., Indrabayu, Suryani, T., Mauludiyanto, A.: A multivariate autoregressive model of rain attenuation on multiple short radio links. IEEE Antennas and Propagation Letters 5, 54–57 (2006)
13. Paulson, K.S., Gibbins, C.J.: Rain models for the prediction of fade durations at millimetre wavelengths. IEE Proceedings - Microwaves, Antennas and Propagation 147(6), 431–436 (2000)
14. Tucker, D.F., Li, X.: Characteristics of warm season precipitating storms in the Arkansas-Red river basin. Journal of Geophysical Research (submitted, 2009)
15. Hou, P., Zhuang, J., Zhang, G.: A rain fading simulation model for broadband wireless access channels in millimeter wavebands. In: Proceedings of the IEEE Vehicular Technology Conference, Tokyo, Japan, May 2000, vol. 3, pp. 2559–2563 (Spring 2000)
16. Crane, R.: Prediction of Attenuation by Rain. IEEE Transactions on Communications 28(9), 1717–1733 (1980)

17. ITU-R P.530: Propagation data and prediction methods required for the design of terrestrial line-of-sight systems. ITU-R Recommendation P.530
18. Folies, J., Raman, B., Jabbar, A., Smith, D., DePardo, D., Sterbenz, J.P., Tucker, D., Euler, T., Frost, V.S.: Experience with frame transfer over a millimeter-wave link. IEEE Communication Letters (submitted, 2008)
19. Waharte, S., Boutaba, R., Iraqi, Y., Ishibashi, B.: Routing protocols in wireless mesh networks: challenges and design considerations. Multimedia Tools Appl. 29(3), 285–303 (2006)
20. Draves, R., Padhye, J., Zill, B.: Comparison of routing metrics for static multi-hop wireless networks. SIGCOMM Computer Communications Review 34(4), 133–144 (2004)
21. Yang, Y., Wang, J., Kravets, R.: Designing routing metrics for mesh networks. In: WiMesh 2005: Proceedings of the IEEE Workshop on Wireless Mesh Networks (2005)
22. Ramachandran, K., Sheriff, I., Belding, E., Almeroth, K.: Routing stability in static wireless mesh networks. In: Proceedings of the eighth Passive and Active Measurement conference, Louvain-la-neuve, Belgium (April 2007)
23. Kim, B.C., Lee, H.S.: Performance comparison of route metrics for wireless mesh networks. IEICE Transactions on Communications 89(11), 3124–3127 (2006)
24. Moy, J.: OSPF version 2. RFC 2328 (Standard) (April 1998)
25. Chandra, M., Roy, A.: Extensions to OSPF to support mobile ad hoc networking. Internet-Draft, draft-ietf-ospf-manet-or-00, Work in progress (February 2008)
26. Spagnolo, P., Henderson, T.: Comparison of proposed OSPF MANET extensions. In: Proceedings of the IEEE Military Communications Conference, Washington, DC, USA, pp. 1–7 (October 2006)
27. Iannone, L., Khalili, R., Salamatian, K., Fdida, S.: Cross-layer routing in wireless mesh networks. In: Proceedings of the 1st International Symposium on Wireless Communication Systems, pp. 319–323 (2004)
28. Pei, G., Spagnolo, P.A., Bae, S., Henderson, T.R., Kim, J.H.: Performance improvements of OSPF MANET extensions: A cross layer approach. In: Proceedings of IEEE Military Communications Conference, Florida, USA, pp. 1–7 (October 2007)
29. Cisco: Cisco IOS IP Command Reference, vol. 2, 4: Routing Protocols. Cisco Systems, Inc. Release 12.3(11)T edn. (2008)
30. Ns-2: The network simulator (July 2008), `http://www.isi.edu/nsnam/ns/`