Survivable Millimeter-Wave Mesh Networks

Abdul Jabbar, Justin P. Rohrer, Victor S. Frost, and James P.G. Sterbenz

Information and Telecommunication Technology Center
Department of Electrical Engineering and Computer Science
The University of Kansas, Lawrence, KS, USA
Email: {jabbar, rohrej, frost, jpgs}@ittc.ku.edu

Abstract

Millimeter-wave mesh networks have the potential to provide cost-effective high-bandwidth solutions to many current bandwidth-constrained networks including cellular backhaul. However, the availability of such networks is severely limited due to their susceptibility to weather, such as precipitation and humidity. In this paper, we present a rigorous approach to survivable millimeter-wave mesh networks based on experimentation, modeling, and simulation. Individual link performance is characterised using frame error-rate measurements from millimeter-wave transmissions on test links over a period of one year. A geometric model based on radar-reflectivity data is used to characterise rain storms and determine their impact on spatially correlated links of a mesh network. To mitigate the impact of link impairments on network services, we present two cross-layered routing protocols to route around the failures: P-WARP (Predictive Weather-Assisted Routing Protocol) and XL-OSPF (Cross-Layered Open Shortest-Path First). We conduct a performance analysis of the proposed mesh network under the presence of actual weather events as recorded by the US National Weather Service. Results show that the proposed approach yields the highest dependability when compared against existing routing methods.

Keywords: millimeter-wave, mesh networks, link measurements, P-WARP, XL-OSPF, predictive routing, cross-layer, storm modeling

1. Introduction

With the ever-increasing number of users and devices that depend on mobile computing, along with the increased sophistication in applications, the demand for high-bandwidth data access to end users is continually increasing. This, in turn,

has led to a significant increase in the capacity requirements for backhaul communication between base stations and the core network, which has traditionally relied on T1 circuits and are currently the bottlenecks in cellular networks. Emerging broadband-access technologies such as 802.16 WiMAX, and LTE (Long Term Evolution) demand backhaul link capacities in excess of a Gb/s. On the other hand, increased deployment of fiber capacity is extremely expensive with more than one third of the operational expenses for wireless operators resulting from backhaul cost [1]. This problem is particularly acute for new 3G service providers that do not have existing fiber infrastructure, and must either lay new fiber or lease capacity from competitors. An alternative that is increasingly attractive for these carriers is to deploy wireless-backhaul links that have significantly lower barriers to entry in both cost and regulation. The need for reliable wireless-mesh technologies to supplement fiber optic networks both in metropolitan and rural areas for practical and economic reasons is well established [2]. Recently, fixed-wireless mesh networks have been emerging to provide cost-effective alternatives for last mile access [3, 4, 5]. There are however, several limitations to existing fixed-wireless technologies. Most notable are the dependability issues due to shared medium, channel susceptibility to impairments such as weather, and long restoration times [6]. Furthermore, the capacity of traditional wireless alternatives is limited, often with rates that are orders of magnitude lower than fiber.

Recently, millimeter-wave links that operate in the 71 – 86 GHz frequency band have been proposed as a cost-effective high-speed alternative for fixed wireless mesh networks [7, 8, 9]. Existing commercial radios in this band can deliver data rates as high as 1 Gb/s with the potential for higher speeds on the order of 10 Gb/s using advanced modulation schemes [10]. Since millimeter-wave radio links require line-of-sight, they are usually deployed with small (30 to 60 cm diameter) high-gain directional antennas resulting in narrow-beamwidth transmissions. These non-interfering pencil beams not only increase the range but also enable spatial reuse, thereby creating mesh-like topologies with increased redundancy and capacity feasible.

High-frequency transmissions are very susceptible to weather disruptions. In particular, millimeter wave transmissions suffer heavy attenuation due to precipitation [7, 11, 12, 13]. As a result, link availability and reliability is significantly impaired during rain storms [7]. Therefore, with these links, it is essential to use mesh topologies, in which there are physically diverse paths between any given node pair. We expect that millimeter-wave mesh networks (MWMNs) will be connected in a arbitrary grid topology as shown in Figure 1, in which several access nodes (e.g A, B, and D) are sending data towards a central office or POP (point of
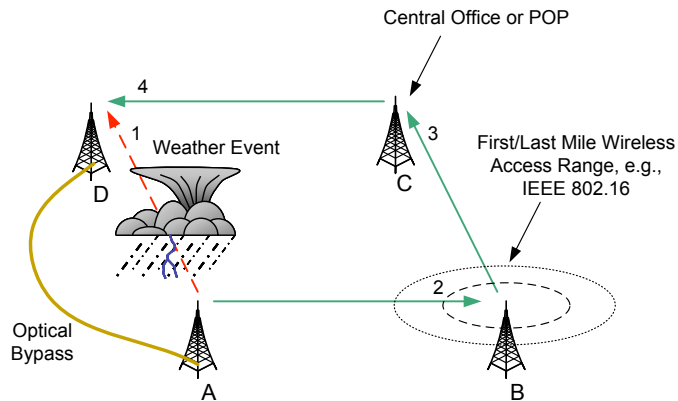
2

Figure 1: Schematic of an MWMN with weather system.

presence, e.g C) that has external access to the wired network. For generality, we show that the nodes may be connected with alternative links (such as fiber or low-frequency radio for redundancy), even though this selection of alternatives is not the focus of this paper. In order to provide dependable paths in such a network, it is essential to design and engineer disruption tolerance into the network. In a mesh topology, this can be achieved through adaptive routing at the network layer.

In this paper, we investigate the effect of weather-related disruptions on deployed 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 based on the link-state algorithm that optimise the network either proactively or predictively. First, we present XL-OSPF, a cross-layered version of the well-known OSPF [14] routing protocol. This proactive approach uses link-cost metrics derived from physical-layer information to maximise the performance of OSPF in the presence of degraded links without the usual penalty of frame loss detection. 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 evaluate the performance of both approaches in terms of efficiency and survivability under weather disruptions. The contributions of this paper are as follows:

3

1. Statistical and temporal analysis of the effects of rain, humidity, and snow on millimeter-wave link performance based on year-round link measurements.
2. A three-state geometric model characterising rain storms using radar reflectivity data.
3. XL-OSPF: a cross-layered proactive routing protocol that uses BER (bit-error rate) estimates to optimise network routes.
4. P-WARP: a predictive weather-assisted routing protocol that routes data around potential link failures *ahead* of time.
5. Performance analysis of millimeter-wave mesh networks and their dependability for end-user applications when subjected to actual rain storms as recorded by US National Weather Service.

The rest of the paper is organised as follows: Section 2 describes related research. Section 3 describes the attenuation effect of precipitation on millimeter-wave transmission and presents measurements conducted on test links for a duration of one year. Section 4 presents a geometric model of rain storms along with the link-stability analysis of a square mesh subjected to several types of rain storms. Section 5 presents the effects of eight observed rain storms, modeled over a metropolitan mesh network. In Section 6, we describe two new disruption-tolerant routing protocols: P-WARP (Predictive Weather-Assisted Routing Protocol) and XL-OSPF (Cross-Layered Open Shortest-Path First). We present a performance comparison of the proposed protocols and the resulting network dependability using simulations in Section 7, followed by conclusions in Section 9.1.

2. Related Work

The impact of rain on millimeter-wave links has been previously studied [15, 16]. The attenuation observed on a millimeter-wave propagation due to rain and other atmospheric effects has been well researched [17, 18, 19, 13, 20]. The ITU-R P.530 recommendation [21] provides a methodology to estimate long term statistics of rain attenuation. This method relies on first determining the precipitation characteristics for a given region (e.g. rain rate exceeded for 0.01% of the year), then calculating the attenuation (in dB/km) using ITU-R P.838 [22], and finally estimating the path attenuation based on the link length. Similar attenuation statistics could be calculated for different precipitation characteristics (e.g. rain rate exceeding 0.001 – 0.1% of the year). Related experimental measurements have been reported for a 3-node mesh network operating at 26.80 – 26.86

4

GHz [7]. In this research, we measure the fame level link characteristics for all rain rates using test links over the duration of a year to precisely determine the impact of precipitation on millimeter-wave links deployed in the area of interest. The objective is to determine the ability of millimeter-wave technology to support the quality of service requirements of the targeted applications such as backhaul under the presence of weather disruptions.

Earlier research on mesh networks (e.g. [6]) has considered several different protocols for routing over wireless topologies including MANET (mobile ad hoc network) protocols such as AODV, DSR, DSDV, and OLSR. However generic MANET protocols do no exploit the inherent geography of a mesh network. Moreover, MANET protocols are optimised for operation in lossy and bandwidth-constrained wireless environments – assumptions that do not hold true in the case of fixed millimeter-wave mesh networks. The focus of more recent routing protocols [23, 24] designed specifically for mesh networks work has been on issues related to the shared medium with resulting contention, mobility [25], and metrics to quantify their effect on data transmission [26, 27, 28]. Recently, there has been an effort to exploit cross-layer information in finding optimal paths in the wireless mesh networks, such as the use of hop-count as the routing metric and use of MAC (medium access control) layer information to overcome congestion for a distance-vector routing algorithm [29]. The predictive wireless routing protocol (PWRP) is based on IEEE 802.11 [30] in which channel measurements are used to provide feedback to routing layer. The recent IEEE 802.11s group is defining a hybrid wireless mesh protocol (HWMP) that combines both reactive and proactive methods to optimise routing [31]. However, the existing schemes are based on shared mediums with bandwidth and mobility constraints. As previously described, MWMNs are neither mobile nor have shared-medium contention issues. Thus, with the exception of the link stability, MWMNs share most of their characteristics with *wired* networks.

Wired routing protocols, on the other hand, assume stable end-to-end paths composed of highly reliable links, in which a link failure triggers routing reconvergence. Millimeter-wave links experience a continually-varying link quality or state, particularly during rain storms. This leads to unstable end-to-end paths over longer durations. Since the commonly used link-state protocol, OSPF (open shortest path first) does not mandate nor 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. Realizing the need for OSPF to handle dynamic network scenarios, several modifications to support mobility are being currently considered [32, 33]. However, in case of

5

static MWMNs, the network dynamics are a result of link disruptions rather than mobility. Lastly, the existing fault-tolerant routing, backup, and restoration mechanisms are not directly applicable because unlike wired links the millimeter wave links have a varying link quality (BER) instead of binary link state.

While the focus of the existing routing protocols is to achieve higher quality of service (throughput and delay) with limited resources, MWMNs with their static, point-to-point, high capacity but unreliable links, present a unique case that can benefit from a domain-specific network solution.

3. Impact of Weather on Millimeter-Wave Links

The existing models such as the ITU-R P.530 recommendation and the Crane model provide statistical distribution of rain attenuation for a terrestrial link. However for the control and performance evaluation of MWMNs, an estimation of real time rain attenuation is needed, which is not provided by these models. Hence, we conducted a field study on a test link for the duration of one year to obtain accurate statistical data on link degradations due to weather events. Furthermore, we also measure the impact of humidity and snow on the millimeter-wave transmissions. In this section, we first present a brief summary of the existing rain-rate attenuation models. We then present our experience with millimeter-wave packet transmission over an extended distance.

3.1. Rain attenuation on millimeter wave transmission

Weather disruption in MWMNs are a direct result of rain attenuation in the millimeter-wave band (71 – 86 GHz). This attenuation is caused by two aspects of rain in the transmission path: scattering by droplets and energy absorption by water. Furthermore, the shape of the rain droplet also affects the signal attenuation. Hence the three major factors that impact the attenuation of a millimeter-wave transmission are signal polarization, center frequency of the signal, and the rain rate.

The most notable models for relating millimeter-wave transmission attenuation to rain rates are the ITU-R P.838 recommendation [22] and the Crane models [16]. The International Telecommunication Union (ITU) outlines a procedure to calculate attenuation from rain rates in ITU-R P.838-3. According to the recommendation, attenuation is given by [22]:

$$\gamma_R = kR^\alpha \tag{1}$$

6

where γ_R is the attenuation in dB/km, R is the rain rate in mm/h, and k and α are obtained from the following equations:

$$\log_{10}k = \sum_{j=1}^{4} a_j \exp\left[-\left(\frac{\log_{10}f - b_j}{c_j}\right)^2\right] + m_k \log_{10}f + c_k \qquad (2)$$

$$\alpha = \sum_{j=1}^{5} a_j \exp\left[-\left(\frac{\log_{10}f - b_j}{c_j}\right)^2\right] + m_\alpha \log_{10}f + c_\alpha \qquad (3)$$

where f is the frequency in GHz, k is either k_{H} or k_{V}, α is either α_{H} or α_{V}. k_{H} and α_{H} are constants for horizontal polarization and k_{V} and α_{V} are constants for vertical polarization. The values for the constants a_j, b_j, c_j, m_k, and c_k are given in [22] for both the vertical and horizontal polarization. For linear and circular polarization, k and α are obtained from the following equations:

$$k = [k_{\mathrm{H}} + k_{\mathrm{V}} + (k_{\mathrm{H}} - k_{\mathrm{V}})\cos^2\theta\cos2\tau]/2 \qquad (4)$$

and

$$\alpha = [k_H\alpha_{\mathrm{H}} + k_V\alpha_{\mathrm{V}} + (k_H\alpha_{\mathrm{H}} - k_V\alpha_{\mathrm{V}})\cos^2\theta\cos2\tau]/2\mathrm{k} \qquad (5)$$

where θ is the path elevation angle ($\theta = 0$ for terrestrial paths) and τ is the polarization tilt angle relative to the horizontal. Equations 2 and 3 were developed from "curve-fitting to power-law coefficients derived from scattering calculations" [22]. In this paper, we use the ITU-R.838 specification to calculate the attenuation γ_R for an observed rain rate of R mm/hr. The total link attenuation is dependent on the length of the link as well as the variation in the rain rate along the link. This attenuation in turn, 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 mesh network is quantified by the bit error rate and subsequently by the frame error rate on individual links.

Next, we discuss the measurement setup used to collect frame error rate measurements over a test millimeter-wave link. We first describe the infrastructure used in this study, followed by a discussion of the impact of rain, snow, and humidity on frame error rate (FER). The link availability results are then presented along with a three-state Markov model.

3.2. Infrastructure Setup

To collect link-performance data, two millimeter wave radio links were deployed in Lawrence, Kansas. In order to observe actual signal attenuation and the

resulting frame errors, we used one radio with no error correction and to evaluate the effectiveness of error correction coding in this specific case, we used a second radio with forward error correction (FEC). As shown in Figure 2, the transmis-
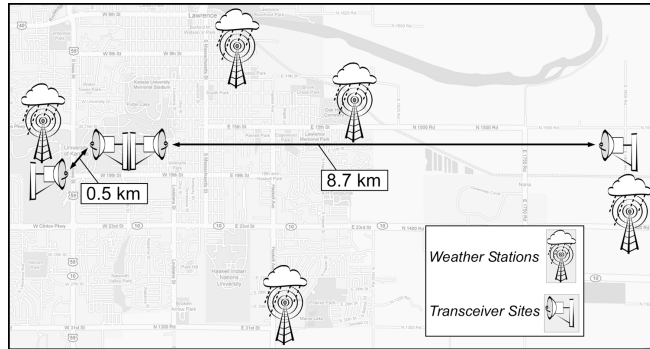

Figure 2: Location of weather stations and radio links

sion path was configured with an intermediate hop; however the primary performance effects arose from the 8.7 km link. To collect weather data for this research project, weather stations using Vaisala WXT510 [34] instruments were deployed at five locations, each equipped with sensors capable of recording temperature, humidity, wind, and most importantly, precipitation [34]. Figure 2 shows the locations of the weather stations and the radio links. Since the purpose of this study is to evaluate real world scenarios, two off-the-shelf radio systems were chosen. Even though the two radios differ slightly in their design parameters, as shown in Table 1, their sensitivity to weather-induced attenuation is similar because of the identical BER vs. SNR curves of the OOK (on-off keying) and BFSK (binary frequency shift keying) modulation schemes. Furthermore, one radio implements a Reed-Solomon (204,188) FEC while the other has no FEC. The objective is to measure the impact of low intensity precipitation on the radios as well as evaluate the effectiveness of FEC in overcoming bit errors.

The link was set up in a loop-back configuration and a GigE tester was used to measure FER along with a number of other performance metrics. To provide input for mesh networking and resilient routing algorithm studies, the link performance metric collected was the FER. Here the FER is the percentage of data frames that are lost during every 30 second interval. The tester transmits on average 3,348,000 512-byte data frames every 30 seconds (0.5 Gb/s). These transmissions, after being looped back at the other end of the line, are recovered at the transmitting end. The FER is then calculated as the percentage of frames lost or in error during

8

Table 1: Radio Parameters

Parameter	Without FEC	With FEC
TX / RX Frequency (FDD)	73.5 GHz /83.5 GHz	72.5 GHz /82.5 GHz
TX Power	17 dBm	17 dBm
Antenna Dia. / Gain / Beamwidth	63 cm / 51 dBi / 0.4°	31 cm / 43 dBi / 0.8°
RX Noise Figure	6 dB	7 dB
Modulation	OOK	BFSK
Data Rate	1250 Mb/s	1000 Mb/s
FEC Type	None	RS (204, 188)

the 30 second interval. Given the total number of frames transmitted, a FER lower than 10^{-7} is not observable. Since it takes 10 seconds to process and record this data, samples were collected every 40 seconds. All observations reported here are from 01 October 2007 through 30 September 2008. Over the course of a single day, around 14,400 (2,880 per station) weather samples and approximately 4,320 (2,160 per link) link performance samples were recorded.

3.3. Effect of Precipitation on FER

The ITU-R P.837 recommendation [35] provides an *estimate* of the precipitation intensities based on *year long* rain rate statistics for a given location in the world. However, due to the significant annual variations, accurate modeling of millimeter-wave links requires precipitation measurements that are time correlated to the link test. As mentioned before, we measured rain intensities at five different locations in and around the link span. During our measurements, we recorded precipitation approximately 1.6% of the time. Figure 3 shows the probability distribution of rain rates for the observed events. The cumulative distribution function for the precipitation is given in Figure 4. It is observed that 95% of the overall precipitation had a rain rate less than 25 mm/hr. The maximum observed rain rate was 160 mm/hr, but occurs with very low probability. As seen in Figure 4, rain rates greater than 60 mm/hr contribute less than 1% towards the overall precipitation. This is important in understanding the characteristics of the rain events, especially that of high-intensity events leading to link disruptions which
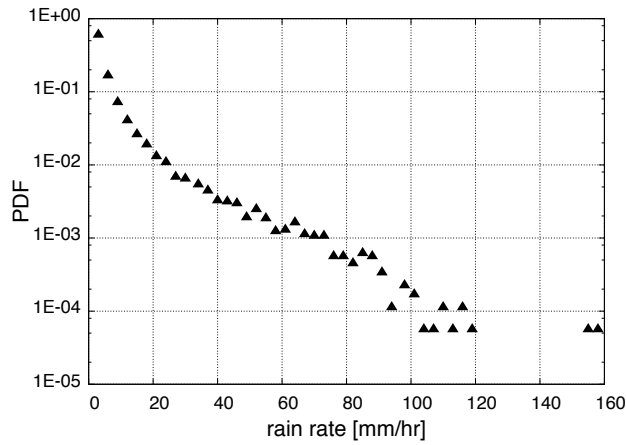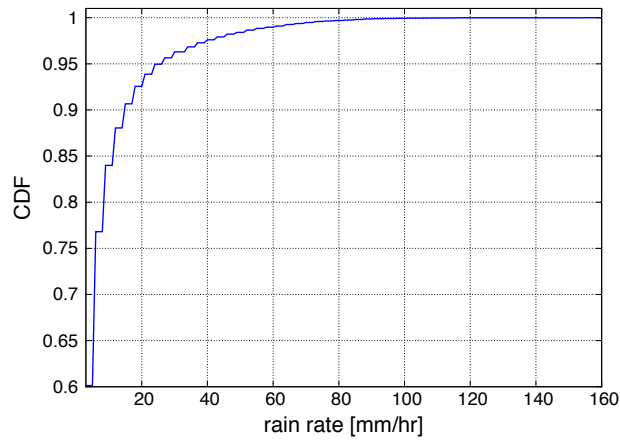
9

Figure 3: Rain rate distribution

Figure 4: Cumulative distribution function of recorded weather events

drive the routing algorithms described in Section 6.

3.3.1. Rain

Figure 5 shows the variation of effective FER during a heavy precipitation event for the system with FEC. It is seen that FEC overcomes degradations to low intensity rain (e.g. below 5 mm/hr). However, for the short-lived, heavy-intensity rain events, the FER can be as high as one leading to complete link failures. For the same event, it was observed that the first radio link, lacking FEC, performed poorly with significant frame errors before and after the rain event due to the latent

Figure 5: FER due to a typical rain event on FEC enabled radio link

moisture in the atmosphere.

3.3.2. Humidity

It is well known that high relative humidity increases the path loss by one dB/km or more depending on the specific conditions [36]. Figure 6 shows variation in humidity as a function of time over a period of five days (03–08 August 2007), together with the associated FER for first radio link. As expected, the FER shows a strong correlation to humidity, tracking its diurnal characteristic. Furthermore, a relative humidity of over 60% is required to observe a noticeable change in the FER. As the humidity increases above 60%, FER increases exponentially. However, an absolute value of 2×10^{-4} FER for a relative humidity of 90% is still small enough to be overcome by Reed-Solomon (204,188) FEC, as implemented in second radio link in which no frame losses were observed.

3.3.3. Snow

Figure 7 shows FER as a function of time for a snow event on 29 January 2008 for first radio link which does not employ FEC. Light to heavy snow was observed between 10:00 and 14:00 hours. As with humidity, the system with FEC did not show sensitivity to snow events, while for the link without FEC the system availability showed significant degradation.

3.4. Availability Analysis

We show the FER distribution on the two test links in Figure 8. Note that these distributions are based on those samples in which FER > 0; in other words, we
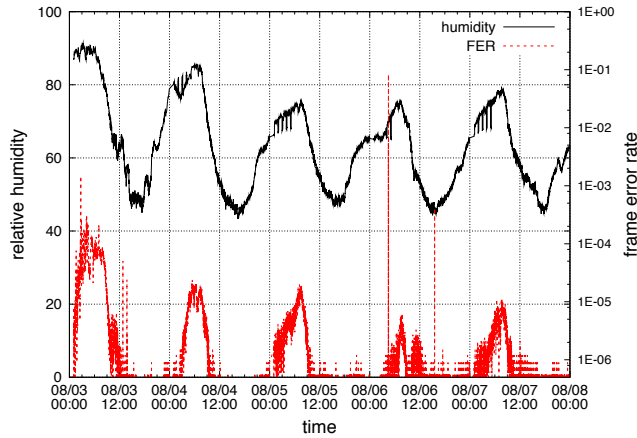
11

Figure 6: Correlation between humidity and FER for the case of no FEC

Figure 7: FER due to a snow event on Jan 29, 2008 for the case of no FEC

calculate the PDF and CDF for *non-zero* error samples. These cumulative distribution functions of the FER shows that the extremes (low and high FERs) contribute more in the error distribution. Comparing the two radios, we see that the CDF of the radio without FEC increases faster in the low FER region when compared to that of radio with FEC. However, the CDF shows that the radio with FEC has fewer (by about 50%) samples with an FER < 1, which means that meaning that when the link with FEC compensates for low error rates and when the channel quality degrades beyond the capability of the FEC, a very high number of sam-
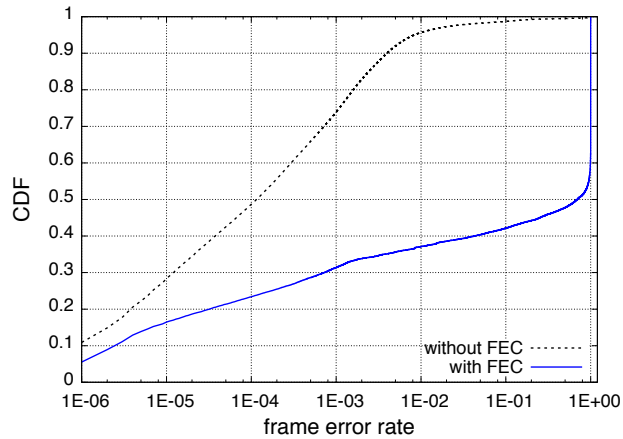
Figure 8: cumulative distribution function of observed FER

ples experience 100% error rate. For a survivability mechanism, this means that there is very small time interval (to adapt) between the time errors begin to appear on the link and the time it fails completely. Of course this drawback is outweighed by the fact that the link with FEC has a much higher average availability, but it is important to note the difference in FER distribution when selecting thresholds for cross-layer routing inputs; if the FER change is gradual parameters should be less sensitive to small changes in FER as compared to the case where FER changes are very drastic. These parameters are further discussed in Section 6.

In this section, we evaluate the availability of the millimeter-wave links based on the data presented above. We define *link availability* as the percentage of time that the link has a FER less than a specified threshold. Each sample of FER is measured over a 40 sec interval transmission as described above. These individual samples are then used to determine link availability over longer time intervals. Availability as a function of FER threshold is given in Figure 9. As expected, these results show the benefit of FEC, however, they also provide an example of the relatively high availability ($\sim$95%) possible for a millimeter-wave link over a long (8.7 km) span.

The availability presented in Figure 9 is specific to the radios used in the experiments. In order to derive a more general measure of link availability under weather disruption, we apply the same analysis to the statistical precipitation data
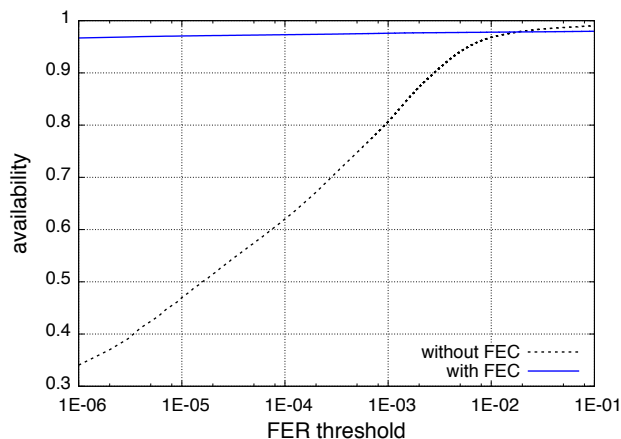
13

Figure 9: Availability as a function of FER threshold

presented in previous section (Section 3.3).[1]

In order to characterise the link from a service perspective, we define three link states: strong, weak and disconnected; each represents a particular range of FER. The thresholds chosen roughly correspond to the current industry standards on the service requirements of cellular backhaul links:

State 1: strongly connected if FER $\leq 5 \times 10^{-5}$

State 2: weakly connected if $5 \times 10^{-5} <$ FER $\leq 5 \times 10^{-2}$

State 3: disconnected if $5 \times 10^{-2} <$ FER

The objective is to quantify link state probabilities based on actual observed weather events (including rain, humidity, and snow) at a given geographical location. Table 2 shows the observed state probabilities for the two radios. These results were used to drive simulation scenarios for the routing protocol evaluation in Section 7.

This shows that, while the RS(204,188) FEC was sufficient to overcome disruptions due to humidity, snow, and low intensity rain, reliable use of millimeter-wave links in the presence of heavy rain warrants the need for solutions above the physical (radio) and link layers.

---

[1]Rain rate statistics vary depending upon the geographic location.

14

Table 2: State probabilities over the duration of the experiments

State	Radio 1 probability (no FEC)	Radio 2 probability (with FEC)
strongly connected	0.3864	0.9469
weakly connected	0.5203	0.0255
disconnected	0.0833	0.0275

4. Modeling Rain Storms

In order to evaluate the effective attenuation on each link in a MWMN, we use the weather-radar echo intensity to determine the overlap of a storm cell with a given link. For the analysis in this paper, we develop a geometric storm model (GSM) to simplify and abstract the calculation of attenuation based on precipitation.

4.1. Geometric Storm Model

In this model, each region in a radar reflectivity map shown in Figure 10 is modeled as an ellipse (R_1–R_5) as shown in Figure 11. While the resolution of the radar reflectivity differentiates many levels of precipitation, the proposed model uses only three levels indicating high (red), low (yellow), and little or no (green) precipitation so that the method remains tractable. The exact value of the rain rate corresponding to a given region depends on the location and the precipitation characteristics. For example, the simulations in our study use a value of 2mm/hr and 5 mm/hr to represent low and high intensity regions respectively. Note that there could be several regions of the same rain intensity in a given storm and the rain rates corresponding to each color vary depending upon the geographical location. Finally, individual links are modeled as line segments. The procedure to calculate the effective attenuation for a given storm as it passes over a fixed mesh network is discussed below.

For a given storm, we first generate the geometrical model of the storm before it enters the geographic region of the network. Secondly, as the storm moves across the network, we generate snap shots of the storm pattern (using ellipsoids) at different time intervals to capture both the regular progress as well as key points when the storm changes direction and the splitting and merging of cells occur[2].

___

[2]This process currently relies to some extent on manual visualization of the radar data. Future
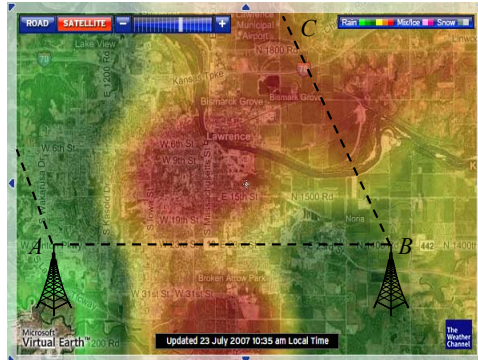
15

Figure 10: Example radar image

Then, with the help of an interpolation program we generated the movement of these ellipsoids over a fixed time step (e.g. one second).

Given that the position of storm ellipses at any time instant t for the entire duration of the simulation, along with coordinates of the nodes and hence the links, we calculate the intersection between the links and the storm ellipses so as to determine the attenuation experienced by each link. Using simple geometry, we calculate the intersection of links with an ellipsoid to determine the length of the link l_i affected by a given storm region i. Using Equation 1, the attenuation of link due to this storm is calculated as $\gamma_R \times l_i$, where γ_R is the rain rate associated with the storm region. The total attenuation on a given link is the summation of the attenuation resulting from all the individual regions that intersect or encompass the line segment[3]. Subsequently, the effective BER is obtained from the attenuation based on the specific radio design. Note that this model considers only the attenuation effects due to rain. While other factors such as humidity contribute significantly lower in the signal attenuation, additional models are needed to capture their impact.

This example shows that while certain links (e.g. $\overline{BC}$) are severely degraded

---

plans include a complete data processing module to automate the process, which has the added benefit of increased rain rate resolution as discussed later in this section.

[3]Note that the quantization of the rain rates in a storm to 3-levels and assumption of constant rain rate inside a region are two simplifying measures in this model. However, the impact of this simplification is conservative estimate of link quality. Future work involves automating this process which would permit the integration of the point attenuation values along the length of the link using the actual value of the rain rate from the radar data as opposed to the quantized values of this 3-level model.
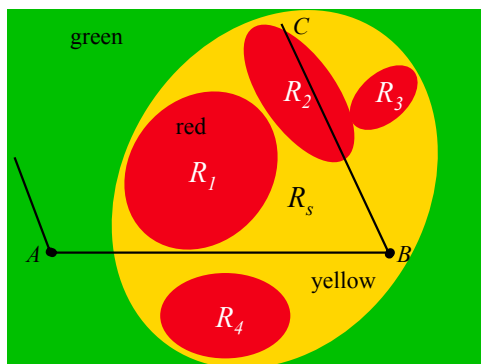
Figure 11: Storm model corresponding to Fig. 10

due to the heavy rain, it is possible to re-route traffic on other adjacent links (e.g. $\overline{BA}$) that are not in an intense rain region. We argue in the following section that this pattern of rain intensity distribution is typical for a majority of storms in the US Great Plains. While Figure 10 shows the storm at one time instant, it can be said without loss of generality that an actual rain event is a series of such snapshots that change over time depending upon the velocity and evolution of the storm. In order to model a storm in real time, we continuously calculate the attenuation for all links as the storm moves through the network. The changes in the link attenuation are dependent on the dynamics of the storm relative to the network[4]

4.2. Radar Measurements in the Midwest United States

In order to determine the impact of actual rain storms, we collected radar reflectivity data from the National Weather Service for the Midwest US. The effect of a weather disruption on millimeter-wave network mainly depends on two factors: rain rate and the geographic footprint of the storm [37]. 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 [38] 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 precipita-

---

[4]The MATLAB code to calculate the attenuation and the resulting BER from a given storm movement is available at http://www.ittc.ku.edu/resilinets/code/ SourceMFiles.zip.

17

tion. We draw two conclusions [37] 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. This is consistent with the our measurements presented in Section 3.3 (See Figure 4). It is important to note that similar studies will be needed for other geographic regions with significantly different weather patterns, such as the Pacific Northwest US and Europe.

In order to get a diverse set of weather patterns in this study, we specifically choose eight storms that were topologically different with significantly different characteristics, e.g. small and large cells, multiple cells on a front line, and intense front line. Their duration over a 1000 km^2 mesh varied from just under an hour to several hours[5]. Figures 12 and 13 each show an instance of two of the eight selected storms.[6,7] These match the previously described common Midwest US characteristic of distinct cells with high intensity.

5. Impact of Weather on Millimeter Wave Networks

In this section we examine the effects of the previously-mentioned eight rain storms moving across a MWMN consisting of a 4×4 grid topology, with 16 nodes and 24 links. The individual link lengths are 10 km and the network spans a region of approximately 1000 km^2 representing a metropolitan-area mesh network. We analyzed the attenuation and BER experienced by all the links in the network during the duration of the storm. The duration of the storm is defined as the time difference between the instance when the first link is affected by the storm and the instant last link recovers from the storm. We then characterise the effect of the individual storms on a per-link basis as well as for the entire network. Finally, we aggregate the results across all storms to get average statistics.[8]

5.1. Channel error rate

As a specific rain event moves across the grid, it affects a number of the links that are in its path. Individual links suffer attenuation to varying degrees depending upon the geographical distribution of the high-intensity regions in the storm.

---

[5]MWMN grid Node 0 is geographically located at 38.8621N, 95.3793W

[6]Distribution 1 observed at 20:39:26 Z on 30 September 2008

[7]Distribution 2 observed at 05:04:11 Z on 22 April 2008

[8]Due to space constraints, we show illustrative results for per-storm analysis from a rain event that was observed on 9 July 2008 in Lawrence, KS, USA

Figure 12: Rain distribution 1 over MWMN links

Figure 13: Rain distribution 2 over MWMN links

The disruptive effect of a given storm on all the MWMN links is presented using a scatter plot of each link BER at every time interval during the storm dura-
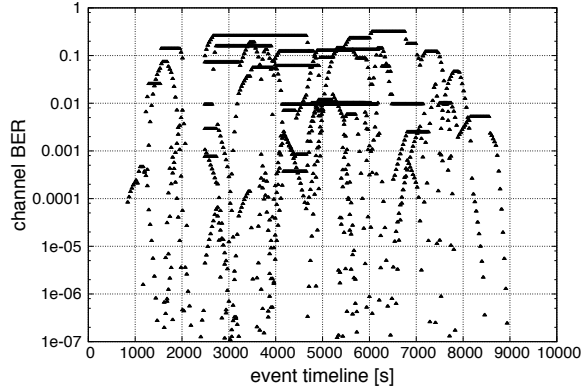
19

Figure 14: Storm effect on link error rates at each time interval

tion, shown in Figure 14. The time interval used for polling was 10 seconds. The plot shows only those links that suffer a BER greater than 1×10^{-7}; all other links were error free. This distribution of BER values indicates that while a number of links were severely degraded, a significant number were either partially degraded or remained normal throughout the duration of the event.

5.2. Mesh availability analysis

As shown in Section 3, the sensitivity of the millimeter wave transmission to precipitation and humidity can be compensated using FEC. Therefore, we calculate the link availability based on the effective BER after the applying the FEC gain from the Reed Solomon (204,188) code. We then quantize the effective BER range into three levels representing the state of the links as *normal*, *partially degraded*, and *severely degraded*. We define *normal* as the state in which the effective BER of a link is less than the threshold of 5×10^{-8}. Links with BER greater than 5×10^{-8} but less than 5×10^{-5} are defined to be *partially degraded*. Finally, links with BER greater than the threshold of 5×10^{-5} are *severely degraded*. These BER thresholds correspond to the FER thresholds mentioned previously. Using these thresholds we determine the percentage of links which fall into a given region and any time during the rain event as shown in Figure 15.

We observe that just before the storm event, 100% of the links in the mesh are in the normal state (state 1), but as the storm moves over the grid a number of links begin transitioning to the partially and severely degraded states. As the storm moves out of the region the links return to the normal state. In order to

20

Figure 15: Percentage of links in each state for each time interval

understand the statistical behavior of the network under these storms we develop a Markov model.

5.3. Three State Model

Based on the thresholds described in section 3.4 we derive a three-state Markov model to characterise the link state. Note that the FER to BER conversion is based on 1000 byte packets. We aggregate the state transitions across all links and all storms.

State 1: normal operation: BER $\leq 5 \times 10^{-8}$

State 2: partially degraded: $5 \times 10^{-8} < \text{BER} \leq 5 \times 10^{-5}$

State 3: severely degraded: BER $> 5 \times 10^{-5}$

Figure 16 shows the state transition diagram, and the transition probabilities are given in Table 3.

6. Cross-Layered Routing

Using the link stability characterization provided in Section 4, we can understand the challenge presented by millimeter-wave mesh networks to mesh routing protocols. Given the topology of Figure 1, the objective of the MWMN routing protocols is to route around link failures without loosing data packets. Sections 6.1 and 6.2 discuss in greater detail two such domain routing protocols originally proposed by the authors in [37] and further developed in [39].

21

Figure 16: Link state transition diagram

Table 3: Transition probability matrix and state probabilities

State	1	2	3	State Probability
1	0.99554	0.00459	0.00046	0.37461
2	0.00417	0.98963	0.01063	0.40449
3	0.00029	0.00578	0.98891	0.22091

6.1. XL-OSPF: Cross-Layered OSPF

A cross-layer approach to link metrics could significantly improve the performance of dynamic link state algorithms [40, 41]. We choose OSPF (Open Shortest-Path First) [14] as the link-state routing protocol due to its wide deployment, use in research, and applicability to fixed networks. OSPF relies on two basic mechanisms to determine link state. One is the link state advertisements (LSAs) generated by each node that carry the status of all its links along with their costs. These 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 updated after the LSAs propagate through the network. With rapidly varying link quality, the only mechanism through which OSPF can detect link degradations is when four consecutive hello packets are dropped, in its default configuration. 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.

The first 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 physical (and MAC) layer which sees the actual packet losses and the network layer which determines the routes. Several mechanisms have been proposed in the literature [25, 29] 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 end hosts of the effective packet error rate. We define a cost metric that is proportional the effective packet error rate. Assuming uniform distribution of the bit errors, the cost of a link between two nodes i and j is calculated as:

$$C_{ij} = P \times \text{BER}_{i,j} \times \gamma \tag{6}$$

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. Further, hysteresis is used with a H_{thresh} 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 primary objective in the MWMN is to avoid disrupted links at all cost, we set this range to 1000. The performance of the modified XL-OSPF with this cost metric is discussed in Section 7.

Even with the error based cost metrics, OSPF requires a finite amount of time before it adapts to changes in link state. If the application or service demands a highly-reliable service, proactive protocols must have a *very* short update interval on the order of milliseconds. But this adds an unacceptable level of overhead, even for broadband networks. In the following section, we discuss a predictive routing scheme that is intended to overcome this problem.

6.2. P-WARP: Predictive Weather Assisted Routing Protocol

As discussed above, proactive algorithms may not be able to meet stringent service requirements (e.g. 50-msec restoration for circuit emulation and CBR traffic) in MWMNs during weather disruptions. Furthermore, it is difficult to measure effective BER or FER at end hosts without an explicit signaling mechanism. In

23

this Section, we investigate the use of information *external* to the network in order to predict the state of links over the next time epoch or several epochs ahead.

The proposed predictive routing algorithm is a link-state algorithm that utilizes weather radar data to forecast the *future* condition of the link. In contrast to the XL-OSPF discussed above, 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 radar reflectivity data modeled in real-time using the methodology discussed in section 4. This processing is done at a either a single *core node* or a small subset of core nodes which are connected to the external Internet and are capable of receiving weather radar data. In either case, multiple-path connectivity into the mesh is necessary for high-availability of the radar data[9]. The topology and physical locations of the (fixed) network nodes are pre-programmed in to the software module that performs the link BER calculation as well as the PER (packet error rate) for a predefined average-packet size. The cost metric for individual links is based on the effective link BER similar to XL-OSPF. While XL-OSPF *proactively* derives costs based on measured BER and propagates updates with conventional LSAs, P-WARP uses short term weather forecast to *predict* link costs. Thus the link cost is calculated using Equation 6 with the same thresholds described in Section 6.1.

6.2.1. Link status updates and route computation

The weather-based link-status *updates* (WLSUs) in P-WARP are slightly different from the conventional 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 receives a WLSU, it recomputes routes using the 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, (and perhaps eliminating) packet loss. It is important to note that while we are using weather predictions to alter the network

---

[9]Note that this connectivity can be provided by low rate links such as lower frequency radio links that are not susceptible to weather

state, the time scale of the weather predictions are on the order of tens of seconds, a short but accurate interval in weather time, however a very long interval in network time, sufficient for predictive routing.

6.2.2. Route sensitivity

It is clearly evident that the effective BER on each link will vary continuously over the duration of the storm. 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. Further, a hysteresis percentage H_{thresh} determines the minimum change in the cost of a link for an update to be generated.

Table 4: Comparison of routing protocols

Metric	Std. OSPF	XL-OSPF	P-WARP
cross-layered	no	yes	yes
link cost C_{ij}	typically 1	$\propto \mathrm{BER}_{ij}$	$\propto \mathrm{BER}_{ij}$
link failure detection	40 s	10 s	0 s
data rerouting	proactive	fast proactive	predictive
control packet	LSA	LSA	WLSU
update rate	periodic 10s	periodic 10s	aperiodic

The various steps in the operation of the P-WARP are enumerated in Algorithm 1. Table 4 shows a brief comparison of the proposed routing protocols XL-OSPF and P-WARP with a standard OSPF implementation.

7. Performance Analysis

In this section, we present the methodology used to model and simulate the storms used in this study. Furthermore, we present the parameters for XL-OSPF and P-WARP used in the simulations. Finally, we quantify the disruptive effect of storms on the MWMN, and compare the disruption tolerance of the proposed mechanisms.

7.1. Simulation Setup

We conducted simulations using storm modeling software that we developed in MATLAB and the ns-2 simulator [42]. The simulated topology consists of 16

Algorithm 1 Predictive Weather-Assisted Routing

Step 1: Generate WLSU at the central node(s)

Input: vertices V, edges E, radar reflectivity data RRD, geographic node positions, forecast window δ_t

1: At time t, receive predicted weather for time epoch $t + \delta_t$
2: update $\leftarrow 0$
3: **for all** $(i, j) \in E$ **do**
4:     $\text{GSM}(t + \delta_t) \leftarrow f(\text{RRD}(t + \delta_t))$ {calculate geometric storm model}
5:     $A_{ij}(t + \delta_t) \leftarrow g(\text{GSM}(t + \delta_t))$ {calculate link attenuation based on link and storm overlap}
6:     $\text{BER}_{ij}(t + \delta_t) \leftarrow h(A_{ij}(t + \delta_t))$ {calculate link BER as a function of its attenuation and radio characteristics}
7:     **if** $\text{BER}_{ij}(t + \delta_t) - \text{BER}_{ij}(t) > \text{BER}_{\text{thresh}}$ **then**
8:        $C_{ij}(t + \delta_t) \leftarrow \text{BER}_{ij} \times P \times \gamma$ {calculate predicted link cost}
9:        **if** $C_{ij}(t + \delta_t) - C_{ij}(t) > H_{\text{thresh}}$ **then**
10:           update $\leftarrow 1$
11:        **end if**
12:     **end if**
13: **end for**
14: **if** update $= 1$ **then**
15:     generate $\text{WLSU}(t + \delta_t)$ consisting of $C_{ij}(t + \delta_t)$, $\forall (i, j) \in E$
16: **end if**

Step 2: Recompute routes at each node

Input: $\text{WLSU}(t + \delta_t)$

1: update local cost matrix C such that $C_{ij} \leftarrow C_{ij}(t + \delta_t)$, $\forall (i, j) \in E$
2: compute least cost paths to obtain $R(t + \delta_t)$
3: schedule $R \leftarrow R(t + \delta_t)$ at time $t = t + \delta_t$

nodes connected in a square mesh as shown in Figure 17. The millimeter-wave links between each pair of nodes are 10 km long. The nodes remain fixed at their locations throughout the simulation.
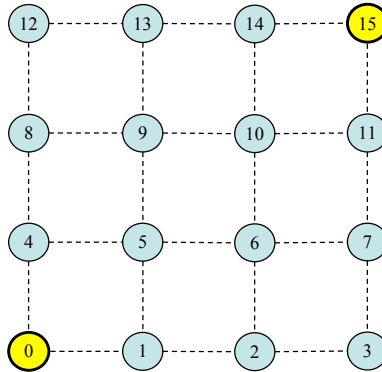

Figure 17: Simulation topology

To model a cellular backhaul network, nodes 0 and 15 are connected to the external network (Internet) and hence all traffic passes through one of these nodes. The remaining 14 nodes generate traffic at a rate of 2.4 Mb/s. The generated traffic is CBR over UDP with a packet size of 1000 bytes. We use eight different storms that took place in the Topeka – Lawrence – Kansas City corridor in 2007 and 2008 to evaluate the disruption tolerance of the proposed mechanisms. These storms, which are modeled using the procedure described in Section 4, consist of an outer ellipse (partially degraded) varying between 30–200 km in diameter and inner ellipses (severely degraded) with a diameter varying between 5–30 km. We have selected two of these storms for illustrative purposes in this paper. The first of these occurred 30 September 2007 in Lawrence, KS, and the second on 22 April 2008. We evaluate the packet delivery ratio and the service availability of the network for four different routing mechanisms.[10]

7.2. Simulation Results

We compare the performance of P-WARP and XL-OSPF to a baseline of conventional OSPF (adapts to link failure but with no cross-layer data input) and static

---

[10]The simulation scripts and input files used are available from `http://www.ittc.ku.edu/resilinets/code/wdtn-2010.tar.gz` and have been run using a standard installation of the ns-allinone-2.33 package.

Figure 18: Windowed average of received packets: first storm

routing that provides a worst-case lower bound on performance. The metrics of interest are packet delivery ratio, delay, and overhead.

7.2.1. Packet Delivery Ratio

The packet delivery ratios averaged over a window of two seconds are shown in Figures 18 and 19 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 the routing protocol. Static routing, as expected, performs very poorly and we show it as a lower performance bound reference. 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 routing 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 error 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 $t = 1400$ sec in Figure 18. At $t = 1420$ sec, the storm disrupts additional links (given that the network is
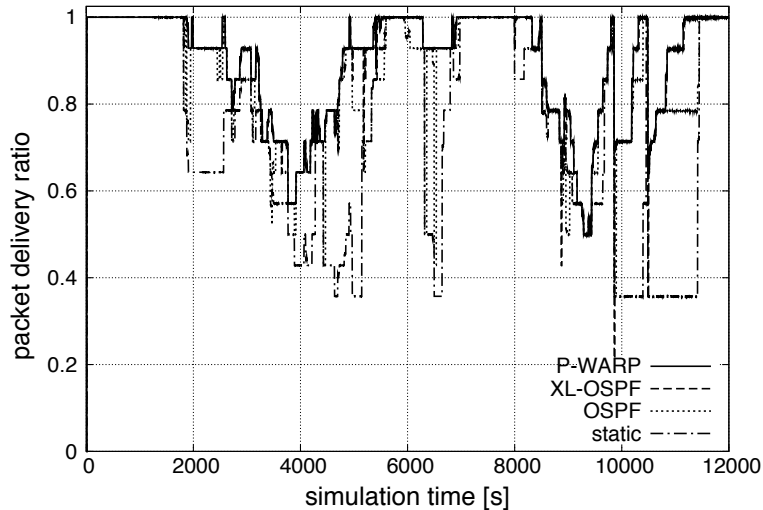
28

Figure 19: Windowed average of received packets: second storm

already degraded to 93% PDR) causing severe packet loss in case of static routing as the PDR drops to 65% and does not recover until the storm has passed at $t = 1800$ sec. On the other hand, OSPF detects failed links and recovers back to the starting PDR at $t = 1460$ sec, indicating a 40s recovery time that corresponds to the dead interval. XL-OSPF recovers much more quickly than OSPF and static routing. It takes approximately 10s (at $t = 1430$) to recover to maximum delivery ratio. Finally, P-WARP maintains the maximum possible delivery ratio indicating a negative reaction (predictive) time. Accurately predicting the impending disruption, P-WARP preemptively routes data on stable paths, thereby avoiding the failed links completely. The reason the maximum possible delivery ratio is not always one is that an intense storm cell located directly on top of a node may affect *all* the outbound links from a particular node causing all packets sourced from and destined to that node to be dropped irrespective of the routing protocol used. The only way to mitigate this effect is to provide alternative paths from a given node: either a fiber connection to the network when practical, or a lower-frequency, lower-bandwidth, but less weather susceptible link such as in the 23 GHz range which also can be used for weather reflectivity data and WLSU dissemination.

In order to compare the aggregate performance of the protocols with respect to each other, the cumulative average of the packet delivery ratio is shown in Fig-
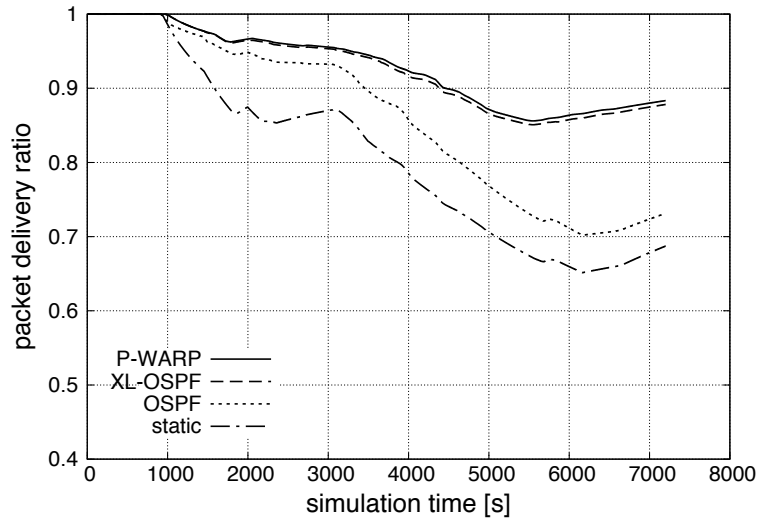
29

Figure 20: Cumulative average of received packets: first storm
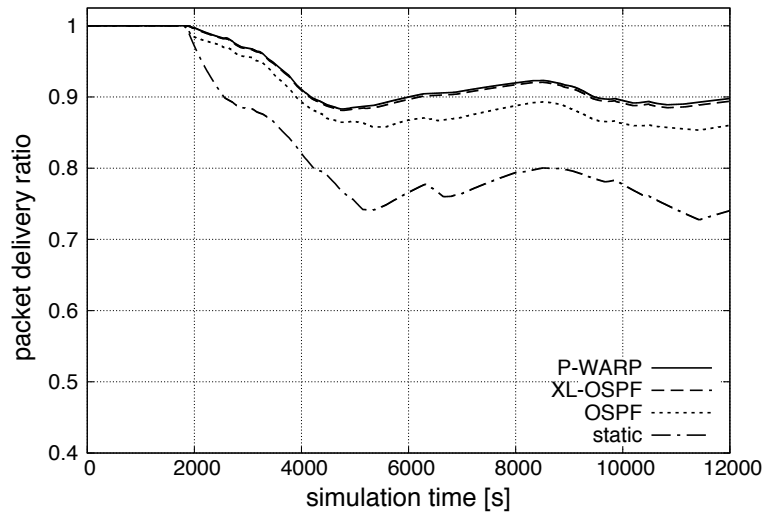

Figure 21: Cumulative average of received packets: second storm
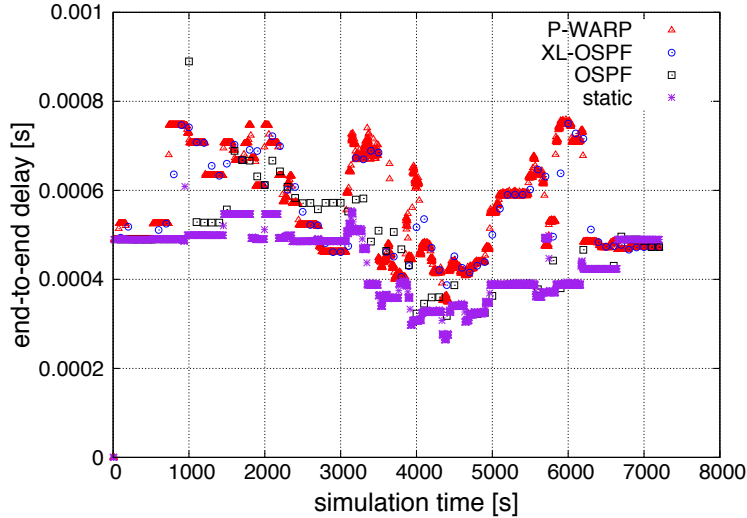
30

Figure 22: End-to-end delay: first storm

ures 20 and 21. The cumulative average of packets delivered by XL-OSPF is very close to that of P-WARP in each case, and both outperform conventional OSPF. In the case of XL-OSPF, the frequency of link state updates determine the reaction time of the network; strict restoration times would require a very small value of update intervals. Because of its predictive nature, P-WARP has two distinct advantages: first, it has *negative* reaction time that might be necessary if stringent service requirements are to be met (such as 50ms restoration time frequently advertised by network service providers); second, the frequency of weather updates does not scale with the restoration times leading to lower protocol overhead.

7.2.2. Delay

While the absolute value of the end-to-end delays in the wireless-mesh topology are negligible, it is worth noting that P-WARP and XL-OSPF can cause higher average delays as they routes packets around link outages, while OSPF and static routing lose more packets during link outages and thus do not incur these delays. Figures 22 and 23 show the end-to-end delays averaged over a two second interval.

We have conducted a number of additional simulations with varying update intervals for OSPF and XL-OSPF. We 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 update frequency, as expected.
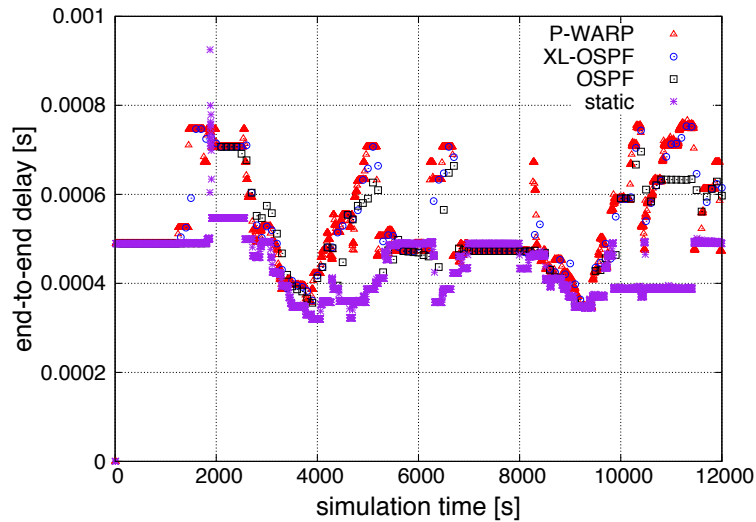
31

Figure 23: End-to-end delay: second storm

7.2.3. Overhead

Static routing does not generate any overhead traffic. Conventional OSPF generates periodic hello messages as well as LSAs. However, the frequency of the LSA does not affect the performance because the quality of the link is not reflected in its cost metric. On the other hand, XL-OSPF uses a cost metric that is proportional to link BER and therefore generates frequent updates that are flooded in the network. As discussed above, in order to react quickly to weather disruptions, XL-OSPF must generate LSAs at a higher rate. This leads to significant increase in the overhead. The number of updates generated in P-WARP are 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. Second, an update is generated only when there is a change in the predicted BER of one or more links. Because of the predictive nature of P-WARP, it is not necessary to generate updates based on the instantaneous changes in the link status *reactively*. Instead, the protocol predicts the cost of all links over a fixed interval and broadcasts these WLSUs to all the nodes. The WLSUs are rate limited to 30 seconds to bound the overhead.

32

8. Applications

In this paper, we have presented an approach to survivable mesh networks based on a combination of fundamentals of millimeter-wave transmission, experimental data collected over test links, modeling of radar reflectivity data, empirical modeling of mesh connectivity, predictive routing protocol, and performance analysis. Consequently, there are a number of assumptions and case-specific parameters that are relevent to our study. In order to apply the same approach to an arbitrary mesh network in an arbitrary geographic location, the steps need to be taken are described below:

1. The bit error rate and attenuation with respect to rain rate, humidity, and snow varies from one radio to another. Hence the sensitivity of the specific radio link should be determined as the first step. This would be either provided by the radio manufacturer or has to be generated experimentally as done in this paper.

2. The weather pattern and the characteristics of the rain storms may vary from one geographic location to the other. Specifically, the size of storm regions with respect to the scale of the link lengths should be determined to evaluate the feasibility of routing around degraded links.

3. If the GSM (geometric storm model) is employed, the mapping of the rate rate to different color coded regions should be evaluated for a specific region.The thresholds that determine the ellipse boundaries should be tuned for the specific case.

4. While the GSM provides an abstraction to reduce the computational complexity of calculation for XL-OSPF and P-WARP, a service provider might chose to perform continuous calculation of attenuation directly on the per-pixel radar images (or data). The tractability and benefit of this approach is left for future research.

5. Lastly, the predicted value of the effective link costs based on the above information should be calculated and fed to the routing algorithm.

9. Conclusions and Future Work

9.1. Conclusions

Millimeter-wave radios may play a significant role in supporting backhaul and other applications for future broadband wireless networks. This paper reports the aggregate frame-level performance of a millimeter-wave link over a one-year

period of observation. The impact of rain, humidity and snow are characterised for a relatively long link based on actual measurements. An analysis of the link availability with respect to FER thresholds and state probabilities were also presented. While it was observed that the RS(204,188) FEC was sufficient to overcome some disruptions due to humidity, snow, and low intensity rain, reliable use of millimeter-wave links in the presence of heavy rain warrants the need for solutions above the physical radio layer.

Furthermore, we quantify the effect of real storms on static millimeter-wave links individually, as well as on the overall state of a millimeter-wave mesh network. Analysis of several observed storms shows that at any given point in time during the event, a small number of links are severely degraded, but there is a large set of links that are either slightly degraded or unaffected. This demonstrates the necessity of a domain-specific routing protocol that uses cross-layer mechanisms to exploit weather predictions and link quality in finding optimal routes. Furthermore, the simulation results show that conventional proactive routing mechanisms do not perform well without a cost metric that reflects the physical state of the link. We compare the performance of two routing protocols that utilise cross-layering between physical and network layer to improve the availability of network. The XL-OSPF protocol uses a cost metric that is proportional to the bit error rate and is shown to perform well under lightly loaded conditions. A predictive routing algorithm, P-WARP, based on short-term weather forecasting, outperforms proactive routing both in terms of throughput and overhead. This algorithm finds the optimal paths with relatively low frequency of routing updates in order to provide reliable service under weather disruptions. Simulations show that millimeter-wave mesh networks with the proposed self-optimising mechanisms form a viable low-cost solution to high-speed backhaul and access.

9.2. Future Work

While this paper presents experimental data on the millimeter-wave link deployment, an empirical model to characterise the impact of weather events remains a part of future work. Preliminary results based on TCP simulations indicate that the performance of the network for transactional traffic is similar to CBR UDP traffic for an over-provisioned network with bursty errors. In the future, we intend to study several more cases to evaluate the performance of the network for different types of applications (e.g FTP, HTTP) as well as different storm scenarios. The simplified geometric storm model quantified radar data in to regions and assumed constant rate in a given region. While this was useful in the initial analysis, future work involves automating this process (via raw radar data processing)

that would permit the integration of the point attenuation values along the length of the link using the actual value of the rain rate from the radar data as opposed to the quantized values of this 3-level model. 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 forwarding decisions. Analysis of a fiber or low frequency radio bypass to overcome node failures remains a part of our future work.

Acknowledgment

[1] Y. Group, Cost optimization of wireless backhaul for next generation migration (March 2005).

[2] S. Ghosh, K. Basu, S. Das, An architecture for next-generation radio access networks, Network, IEEE 19 (5) (2005) 35–42. doi:10.1109/MNET.2005.1509950.

[3] D. Uchida, M. Sugita, I. Toyoda, , T. Atsugi, Mesh-type broadband fixed wireless access system, NTT Technical Review 2 (1) (2004) 44–54.

[4] Y. Wu, J. Hui, H. Sun, Fast restoring gigabit wireless networks using a directional mesh architecture, Computer Communications 26 (2003) 1957–1964.

[5] P. Whitehead, Mesh networks: A new architecture for broadband wireless access systems, in: IEEE Radio and Wireless Conference, Denver, CO, USA, 2000, pp. 43–46.

[6] I. F. Akyildiz, X. Wang, W. Wang, Wireless mesh networks: A survey, Computer Networks 47 (4) (2005) 445–487. doi:10.1016/j.comnet.2004.12.001. URL http://portal.acm.org/citation.cfm?id=1071646

[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] J. A. Khan, H. M. Alnuweiri, 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, 2005, pp. 1–6.

[9] K. Ohata, K. Maruhashi, M. Ito, T. Nishiumi, Millimeter-wave broadband transceivers, NEC Journal of Advanced Technology 2 (3) (2005) 211–216.

[10] E. Torkildson, B. Ananthasubramaniam, U. Madhow, M. Rodwell, Millimeter-wave MIMO: Wireless Links at Optical Speeds, in: Proceedings of the 44th Allerton Conference on Communication, Control and Computing, Monticello, Illinois, USA, 2006.

[11] H. Izadpanah, A millimeter-wave broadband wireless access technology demonstrator for the next-generation internet network reach extension, IEEE Communications Magazine (2001) 140–145.

[12] G. Hendrantoro, Indrabayu, T. Suryani, A. Mauludiyanto, A multivariate autoregressive model of rain attenuation on multiple short radio links, IEEE Antennas and Propagation Letters 5 (2006) 54–57.

[13] K. S. Paulson, C. J. Gibbins, Rain models for the prediction of fade durations at millimetre wavelengths, IEE Proceedings - Microwaves, Antennas and Propagation 147 (6) (2000) 431–436.

[14] J. Moy, OSPF version 2, RFC 2328 (Standard) (Apr. 1998).
URL http://www.ietf.org/rfc/rfc2328.txt

[15] T. Utsunomiya, M. Sekine, Rain attenuation at millimeter and submillimeter wavelengths, International Journal of Infrared and Millimeter Waves 26 (2005) 905–920. doi:doi:10.1007/s10762-005-5662-5.

[16] R. Crane, Prediction of Attenuation by Rain, IEEE Transactions on Communications 28 (9) (1980) 1717–1733.

[17] H. Liebe, An updated model for millimeter wave propagation in moist air, Radio Science, 20 (1985) 1069–1089.

[18] H. Liebe, MPM—An atmospheric millimeter-wave propagation model, International Journal of Infrared and Millimeter Waves 10 (6) (1989) 631–650.

[19] P. Hou, J. Zhuang, G. Zhang, A rain fading simulation model for broadband wireless access channels in millimeter wavebands, in: Proceedings of the IEEE Vehicular Technology Conference - Spring, Vol. 3, Tokyo, Japan, 2000, pp. 2559–2563.

[20] S. Poonam, T. Bandopadhyaya, Rain induced attenuation of millimeter waves radio link in Indian continent, International Journal of Infrared and Millimeter Waves 18 (9) (1997) 1837–1842.

[21] ITU-R P.530, Propagation data and prediction methods required for the design of terrestrial line-of-sight systems, ITU-R Recommendation P.530.

[22] ITU-R P.838, Specific attenuation model for rain for use in prediction methods, ITU-R Recommendation P.838-3 (2005).

[23] S. Waharte, R. Boutaba, Y. Iraqi, B. Ishibashi, Routing protocols in wireless mesh networks: challenges and design considerations, Multimedia Tools Appl. 29 (3) (2006) 285–303. doi:http://dx.doi.org/10.1007/s11042-006-0012-8.

[24] R. Draves, J. Padhye, B. Zill, Comparison of routing metrics for static multi-hop wireless networks, SIGCOMM Computer Communications Review 34 (4) (2004) 133–144. doi:http://doi.acm.org/10.1145/1030194.1015483.

[25] M. Campista, P. Esposito, I. Moraes, L. Costa, O. Duarte, D. Passos, C. de Albuquerque, D. Saade, M. Rubinstein, Routing metrics and protocols for wireless mesh networks, Network, IEEE 22 (1) (2008) 6–12. doi:10.1109/MNET.2008.4435897.

[26] Y. Yang, J. Wang, R. Kravets, Designing routing metrics for mesh networks, in: WiMesh'05: Proceedings of the IEEE Workshop on Wireless Mesh Networks, 2005.

[27] K. Ramachandran, I. Sheriff, E. Belding, K. Almeroth, Routing stability in static wireless mesh networks, in: Proceedings of the eighth Passive and Active Measurement conference, Louvain-la-neuve, Belgium, 2007.
URL http://moment.cs.ucsb.edu/meshnet/datasets/pam.pdf

[28] B. C. Kim, H. S. Lee, Performance comparison of route metrics for wireless mesh networks, IEICE Transactions on Communications 89 (11) (2006) 3124–3127.

[29] H. Q. Vo, Y. Y. Yoon, C. S. Hong, Multi-path routing protocol using cross-layer congestion-awareness in wireless mesh network, in: ICUIMC '08: Proceedings of the 2nd international conference on Ubiquitous information management and communication, ACM, New York, NY, USA, 2008, pp. 486–490. doi:http://doi.acm.org/10.1145/1352793.1352895.

[30] Tropos Networks, Metro-scale mesh networking with tropos metromesh architecture, Whitepaper, Tropos (February 2005).

[31] Status of project IEEE 802.11s: Mesh networking task group report, http://grouper.ieee.org/groups/802/11/Reports/ tgs_update.htm.

[32] M. Chandra, A. Roy, Extensions to ospf to support mobile ad hoc networking, Internet-Draft, draft-ietf-ospf-manet-or-00, work in progress (February 2008).

[33] P. Spagnolo, T. Henderson, Comparison of proposed ospf manet extensions, Military Communications Conference, 2006. MILCOM 2006 (2006) 1–7doi:10.1109/MILCOM.2006.302376.

[34] Vaisala weather transmitter WXT510: The most essential of weather, Vaisala Instruments Catalog (2008).

[35] ITU-R P.837, Characteristics of precipitation for propagation modelling, ITU-R Recommendation P.837-4 (2003).

[36] H. J. Liebe, An updated model for millimeter wave propagation in moist air, Radio Science 20 (1985) 1069–1089.

[37] A. Jabbar, B. Raman, V. S. Frost, J. P. G. Sterbenz, Weather disruption-tolerant self-optimising millimeter mesh networks, in: Proceedings of IWSOS : Third International IFIP/IEEE Workshop on Self-Organizing Systems, Vol. 5343 of Lecture Notes in Computer Science, Springer, 2008, pp. 242–255. doi:http://dx.doi.org/10.1007/978-3-540-92157-8_ 21.

[38] D. F. Tucker, X. Li, Characteristics of warm season precipitating storms in the Arkansas-Red river basin, Journal of Geophysical ResearchSubmitted.

[39] A. Jabbar, J. P. Rohrer, A. Oberthaler, E. K. Çetinkaya, V. S. Frost, J. P. Sterbenz, Performance comparison of weather disruption-tolerant cross-layer routing algorithms, in: INFOCOM'09: The 28th IEEE Conference on Computer Communications, 2009.

[40] L. Iannone, R. Khalili, K. Salamatian, S. Fdida, Cross-layer routing in wireless mesh networks, Wireless Communication Systems, 2004. 1st International Symposium on (2004) 319–323.
URL `http://www-rp.lip6.fr/~iannone/files/IannoneISWCS04.pdf`

[41] G. Pei, P. A. Spagnolo, S. Bae, T. R. Henderson, J. H. Kim, Performance improvements of ospf manet extensions: A cross layer approach, Military Communications Conference, 2007. MILCOM 2007. IEEE (2007) 1–7doi:10.1109/MILCOM.2007.4455134.

[42] ns-2, The network simulator, `http://www.isi.edu/nsnam/ns/` (July 2008).
URL `http://www.isi.edu/nsnam/ns/`