This full text paper was peer reviewed at the direction of IEEE Communications Society subject matter experts for publication in the IEEE INFOCOM 2009 proceedings.

Performance Comparison of Weather Disruption-Tolerant Cross-Layer Routing Algorithms

Abdul Jabbar, Justin P. Rohrer, Andrew Oberthaler, Egemen K. Çetinkaya, Victor Frost, and James P.G. Sterbenz
Information and Telecommunication Technology Center
The University of Kansas, Lawrence, KS, USA
Email: {jabbar, rohrej, ajoj, ekc, frost, jpgs}@ittc.ku.edu

***Abstract*—With growing demand for high-speed access to mobile handheld devices, there is a significant cost benefit in deploying fixed wireless-mesh networks for backhaul access. However, enabling reliable broadband access over high-frequency radios (such as millimeter-wave networks) posses a fundamental challenge due to weather disruptions in general and rain attenuation in particular. In this paper, we present an analysis of the impact of precipitation on millimeter-wave mesh networks based on radar measurements of real storms in the Midwest US. Furthermore, we compare two novel algorithms that use physical-layer information to optimize routing at the network layer: P-WARP (Predictive Weather-Assisted Routing Protocol) and XL-OSPF (Cross-Layered Open Shortest Path First). Finally, we present simulation studies to compare the performance of the proposed protocols and evaluate the dependability of the end-user service during weather disruptions.**

I. 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-speed data access to end users is continually increasing. This, in turn, has led to a significant increase in the capacity requirements for backhaul networks, which has traditionally relied on T1 circuits. 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 is problem is particularly acute for new 3G (and potentially 4G) 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 reliability issues due to shared medium access and channel susceptibility to impairments such as weather. Furthermore, the capacity of traditional wireless alternative is limited, often with rates that are orders of magnitude lower than for fiber. Recently, 802.16 and other line-of-sight wireless solutions are attempting to address some of these issues. More 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 [6], [7], [8]. 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 wireless technologies [9]. Additionally, millimeter-wave mesh networks (MWMNs) can exploit spatial reuse and diversity to increase the degree of node connectivity for redundancy as well as increased capacity. This is possible due to the narrow beam width of the millimeter wave transmissions, which are very narrow non-interfering pencil beams.

High-frequency transmissions are very susceptible to weather disruptions. In particular, millimeter-wave transmissions suffer heavy attenuation due to precipitation [6], [10], [11], [12]. As a result, link availability and reliability is significantly impaired during rain storms [6]. In order to provide dependable paths in such a network, it is essential to engineer disruption tolerance into the mesh network. In a mesh topology, this can be achieved through new routing techniques at the network layer. The objective of this research is to investigate various routing mechanisms that can provide reliable end-to-end paths in the presence of weather disruptions leading to higher service availability.

The contribution of this paper is twofold: first, we present an analysis of the impact of rain storms on millimeter-wave links in a mesh network based on actual radar echo measurements; second, we present a performance analysis of two domain specific routing algorithms (P-WARP and XL-OSPF) that are resilient to link failures. The rest of the paper is organized as follows: Section II describes related research, Section III briefly describes the attenuation effect of precipitation on millimeter-wave transmission followed by link stability analysis of a square mesh subjected to several types of rain storms. In Section IV, we describe the operational details as well as the implementation of two disruption-tolerant routing protocols: P-WARP (Predictive Weather Assisted Routing Protocol) and XL-OSPF (Cross-Layered Open Short Path First) originally proposed in [13]. Finally, we present a performance comparison of the proposed protocols using simulations in Section V.

978-1-4244-3513-5/09/$25.00 ©2009 IEEE

II. Related Work

Earlier research on mesh networks (e.g. [14]) has considered several different protocols for routing over wireless topologies [15], [16]. The focus of previous work has been on issues related to the shared medium (contention), mobility [17], and metrics to quantify their effect on data transmission [18], [19], [20]. 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 the use of MAC layer information to overcome congestion for a distance-vector routing algorithm [21]. The predictive wireless routing protocol (PWRP) is based on IEEE 802.11 [22] 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 optimize routing [23]. However, the existing schemes are based on shared mediums with bandwidth and mobility constraints. As previously described, MWMNs are neither mobile nor have contention issues. Thus, with the exception 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 [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 (bit error rate). Realizing the need for OSPF to handle dynamic network scenarios, several modification to support mobility are being currently reviewed [25], [26]. However, in case of static MWMNs, the network dynamics are a result of link disruptions rather than mobility.

While the focus of the existing routing protocols is to achieve higher quality of service (throughput and delay) with limited bandwidth 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.

III. Impact of Weather on MWMNs

A. 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. If the individual rain droplet shape is more oblate than spherical, the transmission will suffer substantial attenuation if it is horizontally polarized relative to signals that are vertically polarized. 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.530 recommendation [27] and the Crane models [28]. 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 [28]:

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

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

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:

$$\begin{aligned} u &= \frac{\ln(be^{cd})}{d} \\ b &= 2.3 R_p^{-0.17} \\ c &= 0.026 - 0.03 \ln R_p \\ d &= 3.8 - 0.6 \ln R_p. \end{aligned}$$

Based on the above equations, for a given rain rate of R_p mm/hr and a millimeter-wave link of length D, the reduction in received signal strength (RSS) is calculated as A dB. This 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.

B. 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. This is done through a simple geometric model in which each region in a radar reflectivity map shown in Figure 1 is modeled as an ellipse as shown in Figure 2. While the resolution of the radar reflectivity differentiates several 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. Each region is represented as an ellipse (R1-R5) in the geometric model. Note that there could be several regions of the same rain intensity in a given storm. Finally, individual links are modeled as line segments. The attenuation caused on the link due to a single region is then calculated using equation 2 where D is the length of link that is affected by that particular region and R_p is the rain rate of the region. The total attenuation on a given link is the summation of the attenuation resulting from all the individual regions

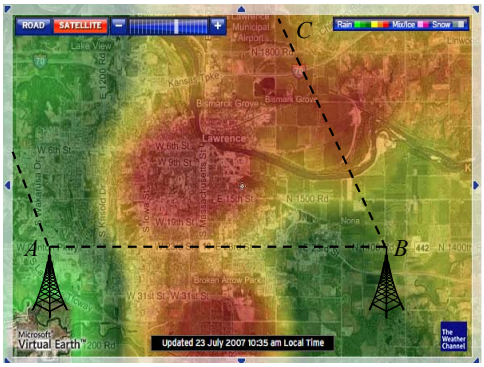

Fig. 1. Example radar image

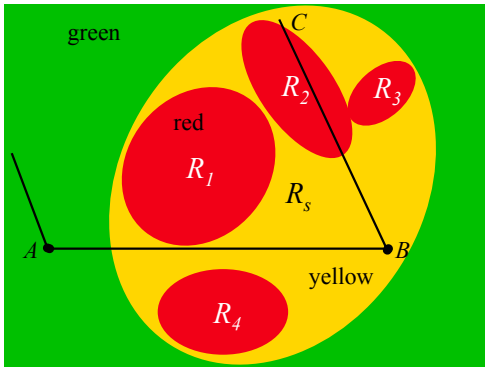

Fig. 2. Storm model corresponding to Fig. 1

that intersect or encompass the line segment. Subsequently, the effective BER is obtained from the attenuation based on the specific radio design.

C. 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 [13]. Both of these parameters vary from one geographic region to the other. For example, based on the analysis of weather data from southern Great Plains of the US [29], authors in [13] draw two conclusions. 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 don't account for a significant percentage of rainfall. 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, we specifically choose eight storms which were topologically different with significantly different characteristics, e.g. small cell, multiple cells on a front line, and intense front line. Their duration over a 1000 km^2 mesh[1] varied from just under an hour to several

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

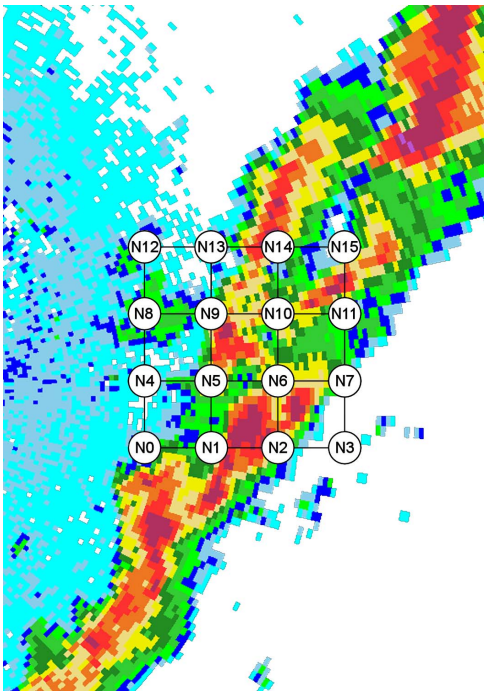

Fig. 3. Rain distribution 1 over MWMN links

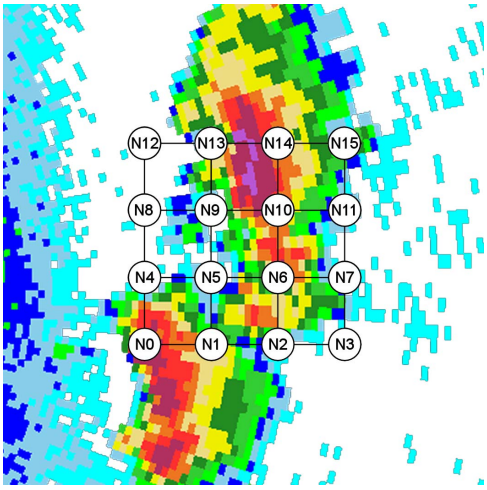

Fig. 4. Rain distribution 2 over MWMN links

hours. Figures 3 and 4 each show an instance of two of the eight selected storms.[2,3] These match the previously described common Midwest US characteristic of distinct cells with high intensity.

D. Link Stability of a Mesh Network

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

[2]Distribution 1 observed at 20:39:26 Z on September 30, 2008
[3]Distribution 2 observed at 05:04:11 Z on April 22, 2008

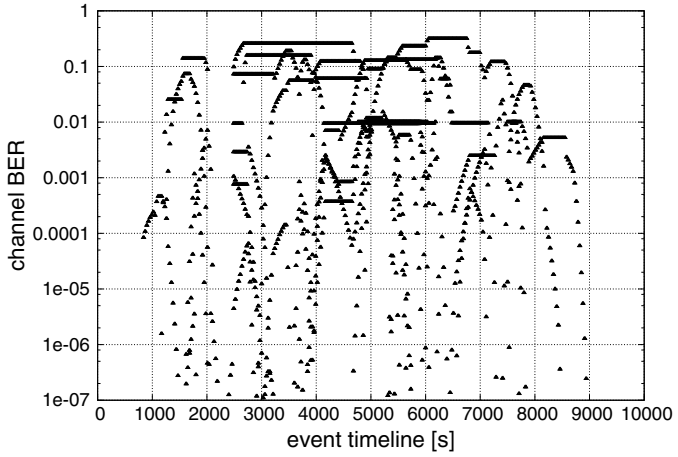

Fig. 5. Storm effect on link error rates at each time interval

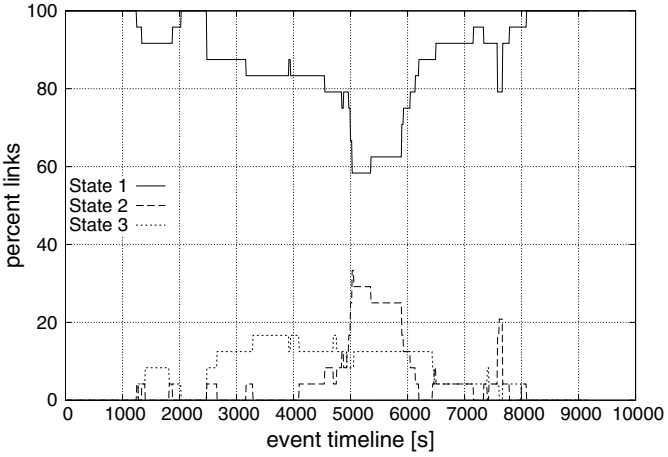

Fig. 6. Percentage of links in each state for each time interval

metropolitan-area mesh network. We analyzed the attenuation and bit-error rate 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 characterize 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.[4]

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.

The disruptive effect of a given storm on the MWMN links is presented using a using a scatter plot of each link BER at every time interval during the storm duration, shown in Figure 5. The time interval used for polling was 10 seconds. The distribution of the 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.

2) Link availability analysis: The sensitivity of the millimeter wave transmission to precipitation and humidity [30] is often compensated using FEC (forward error correction). 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*. 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 6.

We observe that just before the storm event, 100% 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 understand the statistical behavior of the network under these storms we develop a Markov model.

3) Markov model: Based on the thresholds described in section III-D2 we derive a three-state Markov model to characterize the link state. We aggregate the state transitions across all links and all storm.

- State 1 normal operation: BER $\leq 5 \times 10^{-8}$
- State 2 partially degraded: $5 \times 10^{-8} <$ BER $\leq 5 \times 10^{-5}$
- State 3 severely degraded: BER $> 5 \times 10^{-5}$

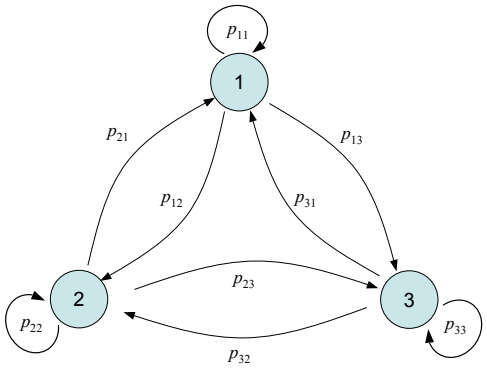

Fig. 7. Link state transition diagram

Figure 7 shows the state transition diagram, and the transitions probabilities are given in Table I.

IV. Cross-Layered Routing

Using the link stability characterization provided in Section III, we can understand the challenge presented by

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

TABLE I
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

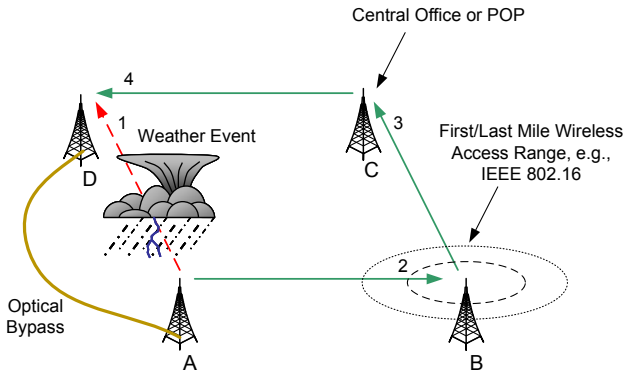

Fig. 8. Schematic of an MWMN with weather system.

millimeter-wave mesh networks to mesh routing protocols.

We expect that MWMNs will be connected in a arbitrary grid topology as shown in Figure 8, in which several access nodes (e.g A, B, and D) are sending data towards a central office or POP (point of 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 (see Section VI-A). Given this topology, the objective of the MWMN routing protocols is to route around link failures without loosing data packets. Sections IV-A and IV-B discuss in greater detail two such domain routing protocols originally proposed by the authors in [13].

A. XL-OSF: Cross-Layered OSPF

A cross-layer approach to link metrics could significantly improve the performance of dynamic link-state algorithms [31], [32]. We choose OSPF (Open Shortest Path First) [24] as the link-state routing protocol for due to its wide deployment and use in research. OSPF has 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 reactively 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 that sees the actual packet losses and the network layer that determines the routes. Several mechanisms have been proposed in the literature [17], [21] 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 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:

$$C_{ij} = P \times \mathrm{BER}_{i,j} \times \gamma \tag{3}$$

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 $\mathrm{BER}_{\mathrm{thresh}}$ of 10^{-8} is used to define the minimum observable change in BER. Further, hysteresis is used with a 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 primary objective in the MWMN is to avoid disrupted links, we set this range to 1000. The performance of the modified XL-OSPF with this cost metric is discussed in Section V.

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, reactive 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.

B. P-WARP: Predictive Weather-Assisted Routing Protocol

As discussed above, reactive algorithms may not be able to meet stringent service requirements in MWMNs during weather disruptions (e.g. 50-msec restoration). Furthermore, it is difficult to measure effective BER or PER (packet error rate) at end hosts without an explicit signaling mechanism. In 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 error packets, we propose P-WARP (predictive weather-assisted routing protocol), in which the BER of each link is calculated from weather radar reflectivity data modeled in real-time using the methodology discussed in section III-B. 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. The topology and physical locations of the (fixed) network nodes are pre-programmed in to the software module that does the link BER calculation as well as the PER 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 reactively derives cost 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 3 with the same thresholds described in Section IV-A.

1) Link status updates and route computation: The weather-based link-status updates (WLSUs) in P-WARP are slightly different from the conventional LSAs (link-state advertisements) 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 state, the time scale of the weather predictions are on the order of tens of seconds, a short but accurate interval in weather time, but a very long interval in network time sufficient for predictive routing.

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. As with XL-OSPF, 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.

V. Performance Evaluation

Recall that the the disruptive effect of real storms on individual link qualities as well as the mesh network as a whole was quantified in section III-D and III-D2, respectively. In this section, we present the simulation of those storms on a typical MWMN network to evaluate the routing performance of the proposed routing algorithms and compare them against existing methods.

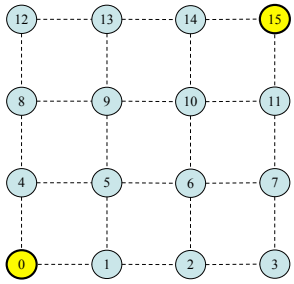

Fig. 9. Simulation topology

A. Simulation Setup

The status of individual MWMN links was extracted from the storm data provided by the National Weather Service, using the geometric model discussed in section III-B. This module was implemented in MATLAB. The resulting link error probabilities were used in ns-2 [33] simulations to evaluate the performance of the routing protocols. The MWMN topology consists of 16 nodes connected in a square mesh as shown in Figure 9. The millimeter-wave links between each pair of nodes are 10 km long. This scenario represent a fixed wireless backbone network with stationary nodes.

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 used 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 III-B, consist of an outer ellipse (severely degraded) varying between 30–200 km in diameter and inner ellipses (partially degraded) with a diameter varying between 5–30 km. Due to space constraints 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. We compare the performance of P-WARP and XL-OSPF to a baseline of conventional OSPF (reactive to link failure but with no radar-data input) and static routing, which provides a worst-case lower bound on performance.

B. Simulation Results

The performance of the routing protocol is measured by its ability to maintain a logical topology that is representative of the current physical topology. When the routing protocol fails to achieve this, packet loss occurs. Therefore, a disruption-tolerant routing protocol must reconverge to new routes with minimum delay in the event of link disruptions. We evaluate this dynamic ability of the protocol by measuring the packet delivery ratio.

1) Packet delivery ratio: The packet delivery ratio for each of the four routing protocols averaged over a window of

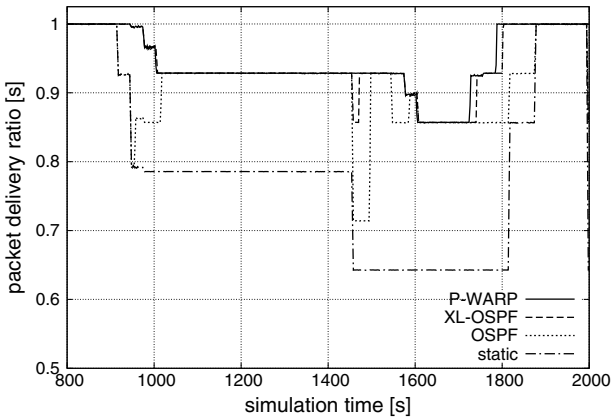

Fig. 10. Windowed average of received packets: first storm

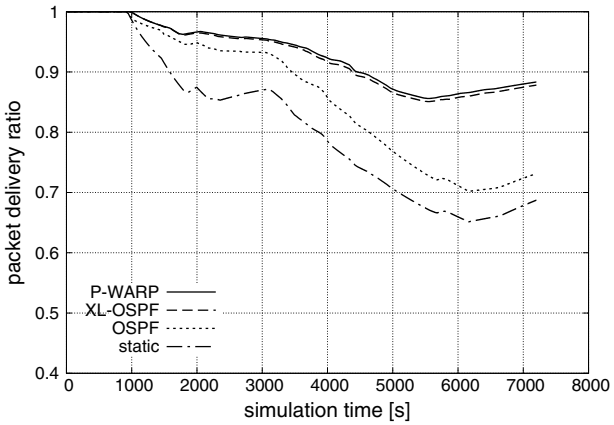

Fig. 12. Cumulative average of received packets: first storm

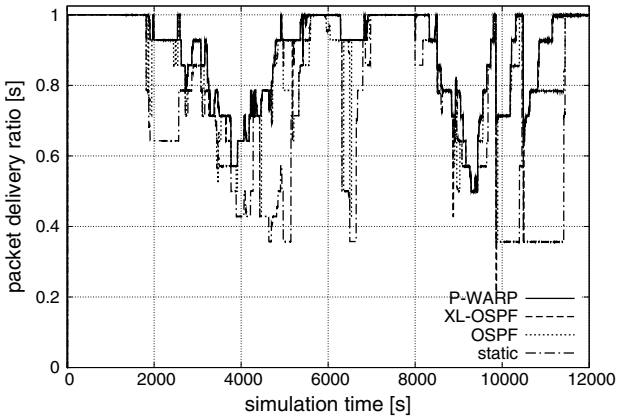

Fig. 11. Windowed average of received packets: second storm

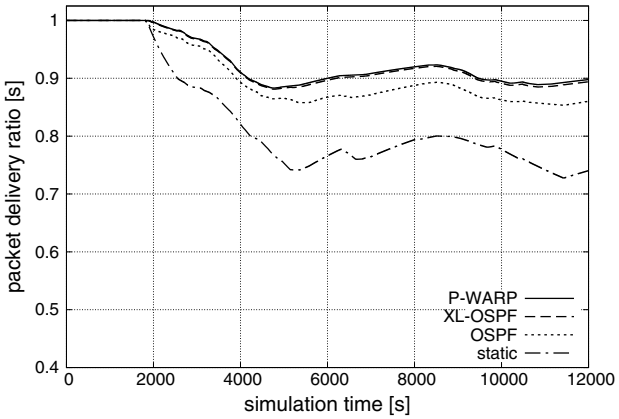

Fig. 13. Cumulative average of received packets: second storm

two seconds is shown in Figures 10 and 11 for two storms illustrated in Figures 3 and 4 respectively. These plots show the instantaneous response of the network to the two storms. As the storm progresses through the network, individual link disruptions result in decreased packet delivery ratio. The extent of this degradation is determined by several factors: the number of links affected, the severity of the link disruption, the existence of alternate routes, and most importantly the ability of the routing protocol to quickly determine optimal routes and reconverge.

The routing protocol used determines the time taken by the network to recover from link failures. As a lower performance bound, we show static routing a 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 and standard OSPF because it can detect degrading links in a shorter time from 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. Note that while the difference between P-WARP and XL-OSPF may seem small in absolute terms, this difference may be important to service providers that have strict service requirements for back-haul links.

For example, consider the packet delivery ratios at $t = 950$ sec. in Figure 10. At 960 sec, the storm disrupts several links causing severe packet loss in case of static routing. Furthermore, the network does not recover until the storm passes at 1875 sec. On the other hand, OSPF detects failed links and recovers at 1000 sec, indicating a 40 sec recovery time that corresponds to the dead interval. XL-OSPF recovers much faster compared to OSPF and static routing. It takes approximately 10 sec 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 1 is that an intense storm cell may affect *all* the outbound links from a particular node causing all packets

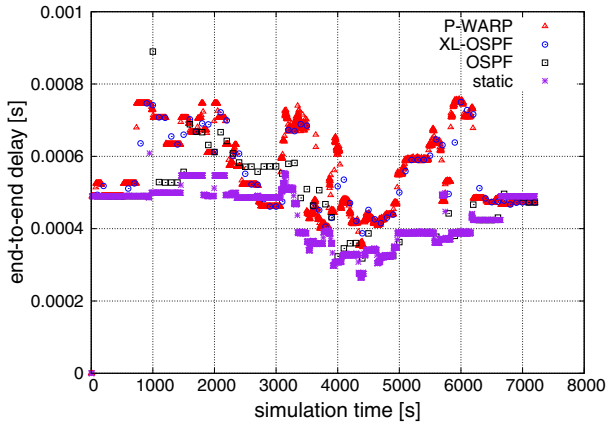

Fig. 14. End-to-end delay: first storm

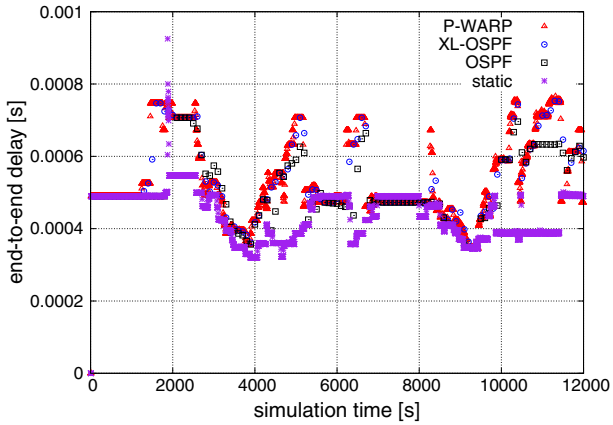

Fig. 15. End-to-end delay: second storm

sourced from 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.

We show the cumulative average of the packet delivery ratio in Figures 12 and 13 in order to compare the *aggregate* performance of the four protocols. The performance of XL-OSPF is very close to that of P-WARP in each case, and they both outperform conventional OSPF. In the case of XL-OSPF, the frequency of link state updates determines 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 (such as 50 msec restoration time frequently advertised by network service providers) are to be met; second, the frequency of weather updates does not scale with the restoration times leading to lower protocol overhead.

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 14 and 15 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.

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 the 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 LS updates at a high rate, leading to a significant increase in the overhead. The number of updates generated in P-WARP are comparatively lower for two reasons. First, a single WLSU generated by the core node(s) carries the predicted costs for all the links. Therefore, link state updates are not generated by individual nodes avoiding the exponential increase in the number of routing messages as the number of nodes increases. Second, an update is generated only when there is a change in the predicted BER of one or more links. Since the performance of P-WARP is not dependent on the arbitrary update frequency, WLSUs are rate limited to 30 seconds to bound the overhead.

VI. Conclusions and Future Work

A. Conclusions

In this paper, 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 reactive 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 utilize cross-layering between physical and network layer to improve the availability of network. The reactive 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 forecast outperforms reactive 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-optimizing mechanisms form a viable low-cost solution to high-speed backhaul and access.

B. 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 type of applications (e.g FTP, HTTP) as well as different types of storms. 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. A fiber or low frequency radio bypass to overcome node failures remains a part of our future work.

ACKNOWLEDGMENTS

We would like to acknowledge the meteorological guidance provided by Dr. Donna F. 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 Bharatwajan Raman. This project was partially funded by Sprint. We acknowledge the technical contributions and guidance of Tim Euler, Technology Strategist, Wireless Backhaul Access Research, Sprint. This work was also partially funded by the NSF FIND (Future Internet Design) program.

REFERENCES

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

[2] S. Ghosh, K. Basu, and S. Das, "An architecture for next-generation radio access networks," *Network, IEEE*, vol. 19, no. 5, pp. 35–42, Sept.-Oct. 2005.

[3] D. Uchida, M. Sugita, I. Toyoda, , and T. Atsugi, "Mesh-type broadband fixed wireless access system," *NTT Technical Review*, vol. 2, no. 1, pp. 44–54, Jan. 2004.

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

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

[6] 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.

[7] J. A. Khan and 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, Sep. 2005, pp. 1–6.

[8] K. Ohata, K. Maruhashi, M. Ito, and T. Nishiumi, "Millimeter-wave broadband transceivers," *NEC Journal of Advanced Technology*, vol. 2, no. 3, pp. 211–216, 2005.

[9] E. Torkildson, B. Ananthasubramaniam, U. Madhow, and 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, Sep. 2006.

[10] H. Izadpanah, "A millimeter-wave broadband wireless access technology demonstrator for the next-generation internet network reach extension," *IEEE Communications Magazine*, pp. 140–145, Sep. 2001.

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

[12] K. S. Paulson and C. J. Gibbins, "Rain models for the prediction of fade durations at millimetre wavelengths," *IEE Proceedings - Microwaves, Antennas and Propagation*, vol. 147, no. 6, pp. 431–436, Dec. 2000.

[13] A. Jabbar, B. Raman, V. S. Frost, and 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*, ser. Lecture Notes in Computer Science, vol. 5343. Springer, 2008, pp. 242–255.

[14] I. F. Akyildiz, X. Wang, and W. Wang, "Wireless mesh networks: A survey," *Computer Networks*, vol. 47, no. 4, pp. 445–487, March 2005. [Online]. Available: http://portal.acm.org/citation.cfm?id=1071646

[15] S. Waharte, R. Boutaba, Y. Iraqi, and B. Ishibashi, "Routing protocols in wireless mesh networks: challenges and design considerations," *Multimedia Tools Appl.*, vol. 29, no. 3, pp. 285–303, 2006.

[16] R. Draves, J. Padhye, and B. Zill, "Comparison of routing metrics for static multi-hop wireless networks," *SIGCOMM Computer Communications Review*, vol. 34, no. 4, pp. 133–144, 2004.

[17] M. Campista, P. Esposito, I. Moraes, L. Costa, O. Duarte, D. Passos, C. de Albuquerque, D. Saade, and M. Rubinstein, "Routing metrics and protocols for wireless mesh networks," *Network, IEEE*, vol. 22, no. 1, pp. 6–12, Jan.-Feb. 2008.

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

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

[20] B. C. Kim and H. S. Lee, "Performance comparison of route metrics for wireless mesh networks," *IEICE Transactions on Communications*, vol. 89, no. 11, pp. 3124–3127, 2006.

[21] H. Q. Vo, Y. Y. Yoon, and 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*. New York, NY, USA: ACM, 2008, pp. 486–490.

[22] Tropos Networks, "Metro-scale mesh networking with tropos metromesh architecture," Tropos, Whitepaper, February 2005.

[23] J. Camp and E. Knightly, "The ieee 802.11s extended service set mesh networking standard," *Communications Magazine, IEEE*, vol. 46, no. 8, pp. 120–126, August 2008.

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

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

[26] P. Spagnolo and T. Henderson, "Comparison of proposed ospf manet extensions," *Military Communications Conference, 2006. MILCOM 2006*, pp. 1–7, Oct. 2006.

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

[28] R. Crane, "Prediction of Attenuation by Rain," *IEEE Transactions on Communications*, vol. 28, no. 9, pp. 1717–1733, 1980.

[29] D. F. Tucker and X. Li, "Characteristics of warm season precipitating storms in the Arkansas-Red river basin," *Journal of Geophysical Research*, 2009, submitted.

[30] H. J. Liebe, "An updated model for millimeter wave propagation in moist air," *Radio Science*, vol. 20, pp. 1069–1089, 1985.

[31] L. Iannone, R. Khalili, K. Salamatian, and S. Fdida, "Cross-layer routing in wireless mesh networks," *Wireless Communication Systems, 2004. 1st International Symposium on*, pp. 319–323, 2004. [Online]. Available: http://www-rp.lip6.fr/~iannone/files/IannoneISWCS04.pdf

[32] G. Pei, P. A. Spagnolo, S. Bae, T. R. Henderson, and J. H. Kim, "Performance improvements of ospf manet extensions: A cross layer approach," *Military Communications Conference, 2007. MILCOM 2007. IEEE*, pp. 1–7, Oct. 2007.

[33] NS-2, "The network simulator," http://www.isi.edu/nsnam/ns/, July 2008. [Online]. Available: http://www.isi.edu/nsnam/ns/