

An Empirical Comparison of Throughput-Maximizing Wireless Mesh Routing Protocols

Rupa Krishnan
Stony Brook University
krishnan@cs.sunysb.edu

Ashish Raniwala
Stony Brook University
raniwala@cs.sunysb.edu

Tzi-cker Chiueh
Stony Brook University
chiueh@cs.sunysb.edu

ABSTRACT

Communication quality of wireless network links is heavily dependent on various external factors such as physical geometry of environmental objects and interference among radio signal sources. As a result, the radio channel quality of real-world wireless networks tends to exhibit both short-term and long-term temporal variations that are in general difficult to model analytically. There has been a large body of research on maximizing the overall throughput of wireless mesh networks through dynamic load/capacity measurement and adaptive routing. However, so far there is no comprehensive evaluation of different protocol mechanisms on a real wireless network testbed. In this paper we first identify the major design dimensions of throughput-maximizing wireless mesh network routing protocols: *wireless link capacity estimation*, *routing path selection*, and *adaptation to temporal link quality fluctuation*, and empirically quantify the performance comparison of various alternatives in each dimension using both software simulations and a miniaturized multi-hop wireless network testbed—MiNT-m.

Categories and Subject Descriptors

C.2.1 [Network Architecture and Design]: Wireless Communication; C.2.2 [Network Protocols]: Routing Protocols

General Terms

Design, Experimentation, Measurement, Performance

Keywords

Wireless Mesh Networks, Routing Protocols, Wireless Testbed, MiNT-m, CARP

1. INTRODUCTION

Radio signal propagation over a wireless link is affected by many factors such as attenuation caused by physical objects, signal strength, interference from other radio signal

sources etc. As a result of the complex interactions among these mechanisms, a wireless link's signal quality and capacity tends to fluctuate significantly over time. Several attempts have been made in the past to model the temporal quality fluctuations of radio channels [20]. However, results of these modeling efforts are largely unsatisfactory for real-world wireless networks [8, 6], because they cannot accurately account for all the effects of external interference, physical obstacles, and low-level radio propagation dynamics. An alternative to analytical modeling is to perform real-time measurements of wireless link quality. This approach incurs a performance overhead, but is more accurate because the result by construction is tailored to a particular network at a particular point in time.

Wireless mesh networks (WMN) are multi-hop wireless networks that are gaining popularity as an alternative last-mile access technology besides DSL and cable modem. Nodes in WMNs are mostly stationary, and a small number of them are connected to the wired network serving as gateways. Since the radio resource available to a WMN is limited, it is essential to maximize the utilization efficiency, even at the expense of increased control and system complexity. In recent years, several throughput-maximizing wireless mesh network routing protocols have been proposed. Some of these such as CARP [13], SRCR [3], MCL [9] are cross-layer in nature and take into consideration the real-time load and capacity of the network's underlying wireless links. In general, the design space of wireless mesh network routing protocols consist of the following dimensions:

- **Wireless link capacity estimation:** how to measure the residual capacity of a wireless link that is shared by multiple nodes in real time.
- **Routing path selection:** how to select the network path with the maximum residual capacity given the residual capacity of individual wireless links, and
- **Adaptation to temporal link quality fluctuation:** how to responsively react to temporal fluctuation in wireless link quality, which should be considered as the norm and designed into wireless routing protocols as a first-class consideration.

Unfortunately previous work do not try to evaluate the various alternatives in each dimension of the design space of wireless mesh routing protocols. The goal of this paper present the results of such a study.

Since these protocols leverage wireless channel information, we carry out this comparative performance study using a multi-hop wireless network testbed called Mint-m [7], which is designed to be reconfigurable in terms of physical

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

WICON '08, November 17-19, 2008, Maui, Hawaii, USA.
Copyright 2008 ACM ICST 978-963-9799-36-3 ...\$5.00.




Figure 1: Wireless link loss variations for 2 test-bed links over a period of 1 hour

network topology/connectivity and link quality, while providing the same fidelity as a real-world multi-hop wireless network testbed. We also used a trace-based simulation approach in which traces collected from the Mint-m testbed are used to drive a software wireless network simulator to simulate large topologies.

2. DESIGN SPACE OF WIRELESS MESH NETWORK ROUTING

First-generation wireless network routing protocols such as AODV [16] and DSR [11] were primarily designed for mobile ad hoc networks, and were designed to maintain network connectivity and reduce the performance overhead of routing protocols. However, most WMN nodes are stationary and typically a few hops from gateway nodes connecting to the wired network. Therefore, the routing protocols for WMNs are more driven to maximize the overall network throughput by incorporating real-time link state and input load information, in a way similar to link-state routing protocols used on wired networks.

Although a WMN's topology is static, the quality of its links varies significantly over time. Therefore the "best" routing path between a pair of nodes also varies over time. As an example, Figure 1 shows the constant and substantial variation in packet delivery ratio of two randomly chosen links on a 12-node Mint-m testbed over a 1-hour period. These measurements were taken at night, when interference from external radio signal sources is minimum, and the testbed's topology remained fixed with minimal human movement that might affect radio channel conditions. The two links were active on orthogonal channels with no interfering traffic. Also, the driver was using the SampleRate algorithm. Similar temporal fluctuations in wireless link quality have been observed by other researchers as well [5, 1]. Therefore, it is essential that WMN routing protocols take explicit account of this time-varying behaviour. In the following subsections we describe the three design dimensions of throughput-maximizing WMN routing protocols in more detail.

2.1 Residual Capacity of a Wireless Link

There are several factor affecting a wireless link's residual capacity. The first factor is the physical transmission rate. The second factor is the interference from radio signal sources in the vicinity of the wireless link. This interference manifests itself in the form of channel contention and bit errors. In addition, previous studies [5] showed that wireless

Metric	Phy Rate	BER	Contention	Impl	Dir
ETX	No	Yes	No	Yes	No
ETT	Yes	Yes	No	Yes	No
PARMA	Yes	Yes	Yes	No	Yes
ICTT	Yes	Yes	Yes	No	No
MTT	Yes	Yes	Yes	Yes	Yes

Table 1: The set of performance factors that existing wireless link capacity metrics incorporate, such as the physical transmission rate (Phy Rate), the bit error rate (BER), the effect of contention from other nodes in the network, practical implementation, and if it is directional (Dir) in nature.

link is directional, i.e. the quality of link $A \rightarrow B$ could be different from that of $B \rightarrow A$. Therefore, the residual capacity of a wireless link should be estimated for each direction separately.

The problem of estimating a wireless link's residual capacity has been extensively studied [6, 10, 22, 23]. In the Expected Transmission Count metric (ETX) [6], the residual capacity of a wireless link is modeled as the expected number of attempts to successfully transmit a packet over the link. To improve over the ETX metric, Draves et al. [10] proposed the Expected Transmission Time (ETT) metric, which incorporates the effects of physical transmission rate and packet loss. The ETT metric of a wireless link is defined as

$$ETT = ETX \times (S/B). \quad (1)$$

where S is the packet size and B is the link's physical transmission rate. One main limitation of ETT is that it does not account for contention from other nodes sharing the channel. To address this problem, the authors suggest a packet-pair technique to estimate the effective value of B . However, this technique requires additional probe packets and therefore can only be applied sparsely. The other drawback of ETT is that it is not directional. The ETT metric is used in both SRCR and MCL protocols.

The Measured Transmission Time (MTT) metric in CARP [13] estimates a wireless link's residual capacity based on *end-to-end per-packet transmission time measurements*. The end-to-end transmission time (TT) of a packet over a wireless link is modeled as

$$TT = T_{access} + N_{retransmit} \times \frac{S}{B} \quad (2)$$

where T_{access} is the time spent in back-off or channel access, $N_{retransmit}$ is the number of retransmissions of the packet, S is the packet's size, and B is the wireless link's physical transmission rate.

MTT approximates it as follows: If there are already packets pending in a wireless interface, then the end-to-end transmission time of the i^{th} packet, TT_i , is equal to the difference between the arrival time of the i^{th} packet's ACK and that of the $i-1^{th}$ packet's ACK. That is,

$$TT_i = Ack_i - Ack_{i-1} \quad (3)$$

where, Ack_i is the time at which the ACK of the i^{th} packet arrives at the sender's interface. If no packet is currently in the wireless interface, then

$$TT_i = Ack_i - Pkt_i \quad (4)$$

where Pkt_i is the time at which the i^{th} packet is moved to the wireless interface. MTT_i is the normalized end-to-end per-packet transmission time, and is equal to $\frac{TT_i}{S_i}$, where S_i is the size of the i^{th} packet in bytes. Finally, a wireless

link's MTT is a smoothed average of MTT_i for all packets transmitted from the link's sender ($MTT_i = MTT_i \times \beta + MTT_{old} \times (1 - \beta)$), where β is a smoothing constant set to 0.25.

In the above formulation, the per-packet queuing time is excluded as it tends to introduce unnecessary fluctuations in the per-packet transmission time measurements as network flows join and leave the network. MTT is a directional metric because the MTT for the $A \rightarrow B$ link is computed independently of MTT of $B \rightarrow A$ link.

The MTT metric has been successfully implemented in the Madwifi driver for Atheros-based wireless LAN interfaces [13]. Several recently proposed wireless link capacity metrics such as MTM [2], PARMA [23] and ICTT [22] also try to incorporate a radio channel's load and/or contention. However none of these metrics have been implemented on real wireless interfaces. Table 1 summarizes the extent to which existing wireless link capacity metrics take into account factors that affect a wireless link's residual capacity.

2.2 Selecting the Maximum-Capacity Routing Path

To select the best routing path between a pair of wireless network nodes, one needs to combine the residual capacity of a candidate path's constituent wireless links, such as MTT values, to derive the path's end-to-end residual capacity. There are two possibilities. One way is to sum up the delay of each link on a path and take the inverse of the resulting end-to-end path delay. The alternative is to take the minimum of the residual capacities of a wireless path's links, and use it as the path's residual capacity.

Because longer routing paths consume more network resources, most existing wireless network routing protocols, for example, SRCR and MCL, take the *summing* approach to favor paths with fewer hops. Although summing the per-link packet transmission times minimizes a flow's end-to-end delay, it does not necessarily maximize a flow's throughput.

We refer to the path capacity metric that is based on the ETT link capacity metric and the summing approach to path capacity estimation as ETTDelay.

$$ETTDelay_p = \sum_{l \in p} ETT_l \times I_l^p \quad (5)$$

CARP supports two path capacity metrics, one approximating the end-to-end path delay and the other approximating the path's available bandwidth.

$$MTTDelay_p = \sum_{l \in p} MTT_l \times I_l^p \quad (6)$$

$$MTTBW_p = \frac{1}{MAX_{l \in p}(MTT_l \times I_l^p)} \quad (7)$$

In all three path capacity metrics, I_l^p represents the degradation of a wireless link l 's capacity because a flow's packets traveling on l also travel on l 's neighboring links on the path p . It represents the degree of *self interference* experienced by a flow traversing the link l on the path p . The value of I_l^p depends on the actual packets being transmitted on interfering links and the physical transmission rates which is hard to measure at a node. Hence we approximate I_l^p by setting it to the number of adjacent links for link l on path p . To illustrate how I_l^p accounts for the effect of self interference, let's consider a wireless path that consists of 3 links, L1, L2 and L3, each of which has a link capacity X .

Mean	Variance (Percent)	TCP Throughput
1.42Mbps	0.71Mbps (50%)	0.019Mbps
0.72Mbps	0.24Mbps (34%)	0.13Mbps
0.475Mbps	0.1Mbps (22%)	0.20Mbps
0.33Mbps	0Mbps (0%)	0.33Mbps

Table 2: The mean and variance of the capacity of four simulated two-link paths, and the throughput of a TCP connection running over each of these paths.

Initially, the path is completely idle, so the initial capacity of the path is $\min(\frac{X}{2}, \frac{X}{3}, \frac{X}{2}) = \frac{X}{3}$. Assume the first flow traversing this path consumes $\frac{X}{6}$ of the path's initial capacity ($\frac{X}{3}$). When the second flow comes along, the measured capacities of L1, L2 and L3 are $\frac{2X}{3}$, $\frac{X}{2}$ and $\frac{2X}{3}$, respectively, and the available path capacity for the second flow is thus $\min(\frac{X}{3}, \frac{X}{6}, \frac{X}{3}) = \frac{X}{6}$, which is exactly the difference between the path's initial capacity and the first flow's consumption.

2.3 Accomodating Temporal Fluctuation of Link Quality

None of the above path capacity metrics account for temporal fluctuations in wireless channel quality. Temporal fluctuations in wireless link quality, as shown in Figure 1, could have a significant performance impact on a WMN's throughput. This subsection presents two methods to address the temporal link quality fluctuation problem.

2.3.1 Modeling Path Capacity as A Probability Distribution

For many transport protocols that are TCP-friendly¹, the *variance* of a wireless network path's capacity is as important as its *mean* because these protocols react very poorly to sudden increase in packet loss. Therefore, in many cases, a wireless path with a smaller mean and variance may well be a better choice from the transport-layer performance standpoint. To demonstrate this, we performed a ns2 simulation for four 2-link paths. The first path goes through two 11Mbps links, each with a loss rate of 50%. The second path goes through a 11Mbps link and a 5.5Mbps, each with a loss rate of 30%. The third path goes through two 5.5Mbps links, each with a loss rate of 50%. The fourth path goes through two 1Mbps links, each with a loss rate of 0%. We measured the throughput of a TCP connection running over each path, and at the same time measured the per-link MTTs, from which we can derive the mean and variance of each path's capacity. The results of this experiment are shown in Table 2. Although the first path has the highest mean path capacity (1.42Mbps), its TCP throughput is the lowest because its variance is the largest (50%). In contrast, the fourth path has the lowest mean path capacity (0.33Mbps), but its effective TCP throughput is actually the highest, because its variance is zero.

Recognizing that both the mean and variance of a wireless path's capacity play an important role, CARP [13] models the capacity of a wireless path p as a probability distribution function unlike other existing WMN routing protocols. To account for path capacity fluctuations in routing CARP computes the complimentary cumulative distribution function (CCDF) $G_l(B)$ of each link's capacity as:

¹A TCP-friendly protocol's sending rate is at most some constant of the square root of the packet loss rate. Also, it not have an unfair advantage over TCP when sharing the same wireless channel.

$$G_l(B) = Pr(x > B) \quad (8)$$

The capacity CCDF of a wireless link (G_{l_i}) is derived from a histogram of the link's capacity measurements over a sliding window.

The CCDF of a path's capacity can be derived from the CCDFs of its constituent links. Assume a N -hop path p consists of the following set of links, l_1, l_2, \dots, l_N , each of whose capacity CCDF is denoted as G_{l_i} . Assume that link capacity distributions are independent, then the CCDF of p 's capacity is expressed as

$$G_p(B) = G_{l_1}(B) \times_{l_2}(B) \times \dots G_{l_N}(B) \quad (9)$$

To choose the best path among a candidate set, CARP computes the maximal capacity that each path is able to guarantee with a probability threshold Th , and chooses the path with the highest guarantee. For example, if Th is set to 0.9, CARP calculates the bandwidth B_{Th} for the i th path such that $B_{Th} = G^{-1}(0.9)$, and chooses the path with the highest B_{Th} . Th is an empirical constant that is set to 0.8 (20th percentile). Mathematically, the best path p_{best} is the one that satisfies the following:

$$p_{best} = arg \max_{p_i} (G_{p_i}^{-1}(0.8)) \quad (10)$$

The above formulation makes no assumptions about the type or form of the distribution functions, and thus could be applied to arbitrary network paths whose capacity CDF is computable. Intuitively, with $Th = 0.8$, CARP chooses the path with the best 20th percentile bandwidth, which is close to a path's mean capacity if it is stable, but may deviate significantly from the mean if the path quality fluctuations are substantial.

CARP assumes that a WMN is connected to the wired backbone through a set of gateway nodes. Every T seconds, each WMN node computes the residual capacity of each outgoing wireless link and sends these link capacity values to its associated gateway node. Each gateway node uses the past N (set to 30) MTT samples to compute the link's capacity CCDF. We will refer to this CDF-based route selection algorithm as *MTTProb*.

2.3.2 Exploiting Diversity via Multi-Path Routing

In real WMNs, over the long term there is no clear winner among a set of routing paths between a pair of WMN nodes, but at any point in time, some paths are better than others. In the case that the best path between a pair of nodes changes faster than the routing protocol can afford to react, the best course of action is to use multiple paths to carry a network flow between the node pair in the hope that they can complement one another at the times when some of them suffer drastic quality degradation.

The two key decisions in the design of a multi-path routing protocol are: (i) which subset of paths to use, and (ii) how to distribute the traffic among these paths. The following algorithm, called *MTTMul*, combines the probability distribution-based path capacity model with multi-path routing, and address these two issues in the following way:

1. Use *MTTBW* metric to compute all possible node-disjoint paths D_p between a pair of nodes, P . At each iteration, add the best path from the remaining network to D_p and remove its links from further consideration. Repeat the above procedure until no more paths are possible. Assume D_p eventually contains K node-disjoint candidate paths.

2. Compute the capacity CDF for each of the K paths found in Step 1. Let the CDF of the i th path be $F_i()$, its offered load be a_i , the end-to-end throughput of the i th path be b_i , and the cumulative distribution function for b_i be $T_i()$. Then

$$T_i(X) = Pr(B_i \leq X) = \begin{cases} F_i(X) & \text{if } X \leq a_i, \\ 1 & \text{if } X > a_i \end{cases} \quad (11)$$

The end-to-end throughput of a path cannot exceed its offered load.

3. The CDF of the total throughput of using all K paths, $T()$, is a convolution of the $T_i()$'s associated with these paths. However in practice, there is inter-path interference among these K paths. Modify the value of I_i^p in Equations 6 and 7 to account for both intra-path and inter-path interference.
4. Let $G_i() = 1 - T_i()$. To search for the best load distribution a_i 's that maximizes $G^{-1}(0.8)$, initially distribute the traffic load among the K paths proportional to $G_i^{-1}(0.8)$. Randomly select a number of traffic distributions, repeat Step 2 and 3 for each of them, and choose the one that maximizes $G^{-1}(0.8)$.
5. Compare the 20th percentile bandwidth of the best multi-path routing choice with that of the best single-path routing choice from *MTTProb*, and choose the better of the two. This way multi-path routing is used only when it is advantageous to do so.

3. PERFORMANCE ANALYSIS AND COMPARISON

3.1 Evaluation Platform

We evaluated the performance of different routing schemes on the MiNT-m testbed [7, 8], which is a 12-node multi-hop wireless network testbed developed at Stony Brook University. Each physical MiNT-m node maps to a wireless node in each experiment and runs an instance of the hybrid ns2 simulator [7, 8], whose link layer encapsulates each ns2 packet into a UDP payload and transmits it to the destination through its physical WLAN interface. Since packets are transmitted through actual software drivers and wireless interfaces over a radio channel, the MAC/physical layer emulation fidelity of MiNT-m is as good as a real-world wireless testbed. In the following experiments, we assumed every WMN node has only one NIC and they all operate in the same channel.

Additionally, for large simulated topologies, we took a trace-driven approach. More specifically, we collected link-layer packet transmission traces from three different types of links on the Mint-m testbed. The three types of links correspond to good (< 5% loss rate), medium (5 – 40% loss rate), and poor (> 40% loss rate). For each link type, we collected multiple packet transmission traces. and played them inside the ns2 simulator. As a result, the packet loss behavior of a simulated link is similar to a testbed link in terms of both average packet loss rate and temporal burstiness of packet losses. However, this approach does not improve the ns2 interference model.

3.2 Evaluation Methodology

We used a 3x4 grid topology for evaluation, and set the WLAN cards to operate in the 802.11a mode. All the nodes operated on the same channel. A node is designated the gateway, and once every second each MiNT-m node sends a

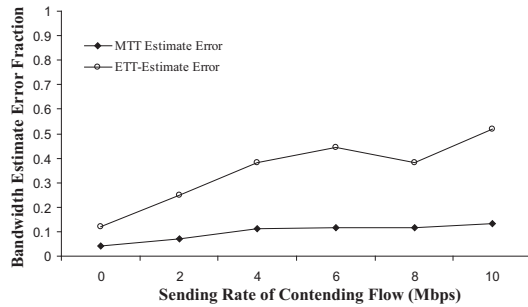


Figure 2: The ratio of error in estimated bandwidth to the actual bandwidth for ETT and MTT metrics. Actual bandwidth is the application level throughput for the one-hop UDP flow. The background traffic is another contending UDP flow that varies its sending rate from 0Mbps to 10Mbps. ETT’s estimation error increases with background load, while MTT’s estimation error almost remains constant.

link state update to the gateway node. In each run, the gateway first computes a path between itself and every testbed node based on link quality information of the past 30 seconds under a background traffic load, using the different routing schemes described previously, then a set of test connections runs for 20 seconds using the computed paths and their throughputs are measured, analyzed and reported. We used the modified Madwifi driver for the MTT metric and provided the following measurements to user-level programs: average MTT, physical transmission rate and retransmission count for each outgoing link.

3.3 Results and Analysis

The goal of this performance evaluation study is to empirically compare the performance of different routing schemes, and to characterize the performance contribution of component mechanisms. The main performance metric used is connection throughput expressed in bits/sec. We compare the MTT link metric with the ETT link metric and compare MTT path metrics with *ETTDelay*. Note that in [10], the authors use a packet-pair based technique to determine B , which requires explicit probing of each link with a packet pair. This can be expensive especially in a loaded channel and would not accurately capture short term channel variations or auto transmission rate adjustment. Hence, for our evaluation we have set the value of B to the physical transmission rate reported by the driver. Also, the authors used the average delivery rate along the link based on a set of probe packets to compute ETX of the link. We modified it to use the average Retry Count (RX) provided by the card. RX has been shown in [5] to be a better metric than ETX. The modified ETT metric was also directional.

$$ETT_{A-B} = RX_{A-B} \times S/B \quad (12)$$

3.3.1 Accuracy of Link Metrics

To evaluate the effectiveness of different wireless link capacity metrics, we compare the application-level throughput of a UDP flow under different background loads to the predictions from MTT and ETT metrics measured at the source node of the UDP flow. The UDP flow sends at 6Mbps over an IEEE 802.11a link that is configured to physically transmit at 6Mbps. This ensures that the link under consider-

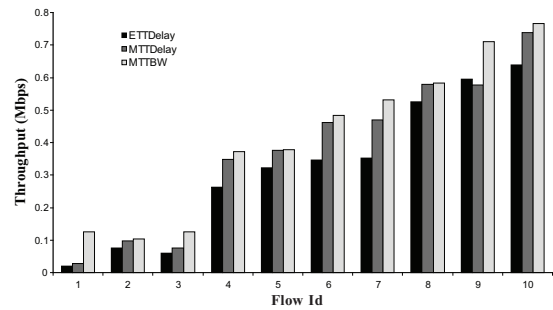


Figure 3: End-to-end throughput measurements of 10 best flows on the MiNT-m testbed for ETTDelay, MTTDelay and MTTBW schemes in the presence of a background flow with a sending rate of 2Mbps. MTTDelay and MTTBW always outperform ETTDelay. MTTBW performs the same or better than MTTDelay in all cases.

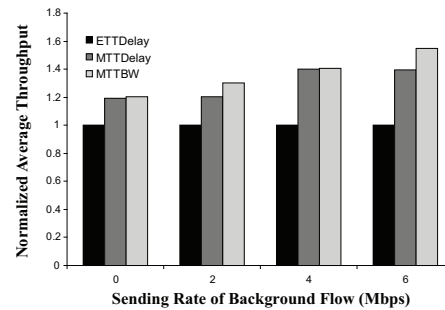


Figure 4: Normalized average throughput of the 20 flows on the MiNT-m testbed for ETTDelay, MTTDelay and MTTBW when the background traffic load is varied from 0 Mbps to 6 Mbps. MTTDelay and MTTBW are more effective than ETTDelay because it does not explicitly take into account contending traffic. MTTBW shows 20-60% improvement over ETTDelay and upto 15% over MTTDelay

ation is saturating the wireless channel in the absence of interference. The background load is in the form of another UDP flow over a neighboring IEEE 802.11a link that operates in the same channel but at the physical transmission rate of 24Mbps.

Figure 2 shows the estimation errors of MTT and ETT as the background UDP flow’s sending rate varies from 0 to 10 Mbps. As the background load of the radio channel increases, the throughput estimation error of ETT increases since it does not consider contention and interference from the background source. In contrast, MTT’s estimation error remains largely the same even in the presence of large background loads. Although MTT accounts for all delay components associated with a packet’s transmission, it does not account for all performance overheads and therefore still exhibits a 5-15% estimation error.

3.3.2 Comparison among ETTDelay, MTTDelay and MTTBW

We next investigate how different path metrics compare with each other. To compare them, we compare the paths selected based on these three path capacity metrics. Because our testbed is a small one, we focus on 20 randomly chosen node pairs. In each run, one of the nodes is designated the

BG Load	ETTDelay	MTTDelay	MTTBW
0Mb	2.72	2.17	2.25
2Mb	2.70	2.23	2.27
4Mb	2.68	2.34	2.34
6Mb	2.72	2.44	2.43

Table 3: The average path length of 20 flows (from figure 4), as the sending rate of the background flow increases. The average path length for ETTDelay does not change as it does not adapt to the increasing background load, whereas the average path length for MTTBW and MTTDelay increases slowly as flows are routed through more hops to stay away from congested links.

gateway, and for the first 30 seconds all nodes send link updates to this gateway. At the end of the 30 seconds, the non-gateway node of each pair sends a route request to the gateway, which computes a route and sends back a reply, and the non-gateway node starts a test flow after receiving the gateway’s reply. The runs are non concurrent, and we repeat this for all the three metrics consecutively to ensure fair comparison. To ensure that we have sufficient interfering traffic we use a 1-hop flow in the center of the testbed as background traffic. This background flow has a sending rate of 2Mbps, and it is present throughout the experiment.

The average throughputs of the 10 best flows under the three routing metrics are shown in Figure 3. In almost all cases, MTTDelay and MTTBW metrics outperform ETTDelay, thus further validating that MTT is a better capacity metric than ETT. In addition, MTTBW performs equally well or better than MTTDelay. This suggests that as far as connection throughput is concerned the “bottleneck bandwidth” approach is more effective than the “sum of delay” approach, and that paths with higher bandwidth could incur longer delay in many cases.

To assess the impact of the background load’s volume, we repeated the same experiment but varying the sending rate of the background traffic from 0Mbps to 6Mbps. This is to show the effect of increasing background load on the path capacity metrics. One problem with high background traffic load is that it increases the bit error rate and in turn causes route request packets to be dropped. In a real wireless network, this is not a major problem because lost packets are retransmitted. However this problem seriously hampers experiment repeatability. Therefore we cap the background traffic load to 6Mbps, lowest rate which can saturate an 802.11a channel.

Figure 4 shows the average throughput of the 20 flows for each path capacity metric. The throughput values are normalized with respect to ETTDelay. To ensure fair comparison, throughput measurements for different path capacity metrics for a particular node pair and background load were run back to back. As expected, the performance gains of MTT-based metrics over ETTDelay increase with the background load, because the difference between ETT and MTT is more likely to manifest itself under heavier loads.

To analyze further, we show the average path length of the 20 flows under the three path capacity metrics in Table 3. The average path length for ETTDelay remains unchanged as the background load increases, as this path capacity metric is less sensitive to a particular constituent link’s load. In contrast, the average path lengths for MTTDelay and MTTBW increase with background load as they steer away from congested links by using longer paths. Note that the

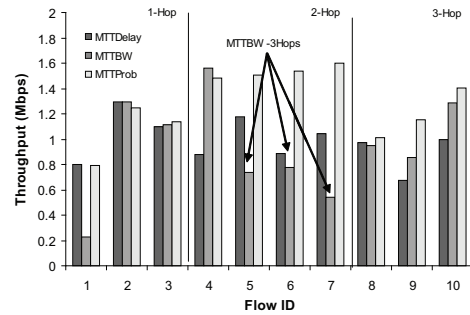


Figure 5: Throughput of 10 TCP Flows on the MiNT-m testbed for MTTProb (T=0.8), MTTDelay and MTTBW. Flows are classified by the number of hops in the path chosen by MTTProb. MTTDelay outperforms MTTBW for Flow 1,5,6,7 because MTTBW chooses a longer path for those flows.

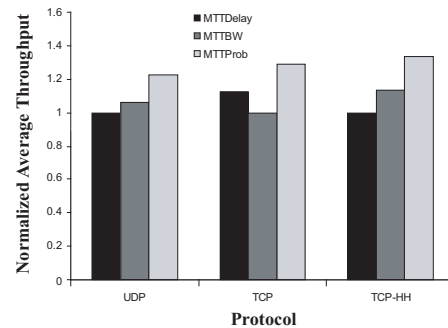


Figure 6: Normalized average throughput of 20 flows on the MiNT-m testbed under three routing metrics, MTTProb, MTTDelay and MTTBW and three transport protocols, UDP, TCP and TCP-HH. MTTProb improves over MTTDelay by 20% for UDP, by 28% for TCP and 33% for TCP-HH

paths ETTDelay chose are generally longer than those chosen by MTTDelay and MTTBW in this particular setup, because ETTDelay used the physical transmission rate for the value of B (Equation 1), and thus was more likely to select paths that go through nodes with higher physical transmission rates, even if the links experience lower throughput due to interference.

3.3.3 Path Capacity as A Probability Distribution

We have established that MTTDelay and MTTBW perform better than the ETTDelay metric in terms of throughput. To compare MTTProb, which approximates a wireless path’s capacity as a probability distribution, with MTTBW and MTTDelay, which represent a wireless path’s average capacity, we selected two node pairs in the MiNT-m testbed, and for each pair picked two alternative paths. Let P_{11} and P_{12} be the two alternative paths for the first node pair, and P_{21} and P_{22} for the second node pair. For each of these paths, we collected the link quality characteristics for 30 seconds and used the resulting information to derive the mean, variance and path’s capacity with particular probability. Then we measured the actual throughput of each path and used it as the ground truth for comparing routing metrics. The results of these measurements are shown in Table 4, where T is set to 0.7, 0.8 and 0.9. The *mean*

Path	Mean	Var	Y=0.7	Y=0.8	Y=0.9	Throughput
P_{11}	0.679	0.134	0.65	0.60	0.50	0.695
P_{12}	0.682	0.260	0.65	0.55	0.45	0.470
P_{21}	1.10	0.310	1.0	1.0	1.00	0.829
P_{22}	1.31	0.453	1.2	1.1	0.95	0.983

Table 4: The mean, variance and path capacity with probability Y and the actual measured throughput of four paths over the MiNT- m testbed. P_{11} and P_{12} are for one node pair, and P_{21} and P_{22} are for another node pair.

capacity alone is not a reliable metric, because P_{12} is worse than P_{11} in terms of actual throughput even though P_{12} has a higher mean than P_{11} . In this case, the main cause is the high variance. $T = 0.8$ seems to be the best choice in this case as routing decisions based on it are consistent with those derived from actual measured throughputs.

In general, higher T results in a more conservative path capacity estimate whereas lower T tends to over-estimate the actual path capacity. How T can be adapted to a given transport protocol such as TCP according to the impact of link quality fluctuation on its throughput is worth exploring further. Because the path capacity CDFs are computed from histograms which contains discrete values, using a percentile path capacity estimate could often lead to a tie among multiple candidate paths. Currently we resolve ties by choosing the path with the smaller fractional variance, i.e., $\frac{\text{variance}}{\text{mean}}$.

A probability distribution-based path capacity estimate is particularly useful for transport protocols such as TCP that are sensitive to fluctuations in path quality. Due to TCPs sensitivity to loss rate the testbed topology was carefully setup to have lower bit error rates 10% for most links and at worst 30%. Setting up such a topology requires positioning nodes such that link quality to most neighbours is good, as opposed to the previous experiment where ensuring that the network is connected was sufficient. We also had to choose the background traffic carefully. A high rate background flow can increase the loss rate of the channel and significantly affect TCP performance. (The performance limitations of TCP on wireless are well known [17, 19]). This would also make repeatability of experiments across different routing schemes impossible. Therefore, the background traffic consisted of 3 UDP flows placed at different points in the grid, each sending traffic at the rate of 200Kbps.

We ran 20 TCP flows over 20 randomly selected node pairs on the MiNT- m testbed. Here again, one of each node pair plays the role of a gateway. The experiments were repeated 3 times for each scheme and there were 180 runs in total. The measured throughputs of best 10 of these 20 TCP flows under MTTProb ($T=0.8$), MTTDelay, MTTBW are shown in Figure 5. Except for two flows (2 and 4), MTTProb performs better than the other two, because it prefers stable paths with lower variance. Also, the average throughput for MTTBW is lower than that of MTTDelay. This is because MTTBW sometimes chooses longer paths with higher mean capacity and a larger number of hops tends to increase the packet loss probability. This adversely affects TCP because it backs off unnecessarily whenever it detects consecutive packet losses.

Because TCP is not the most effective transport protocol for multi-hop wireless networks, we repeated the same experiment using both hop-by-hop TCP (TCP-HH) and UDP. Figure 6 shows the normalized average throughputs of the 20 flows under three routing metrics, MTTProb ($T=0.8$), MTTDelay and MTTBW, for three transport protocols, TCP,

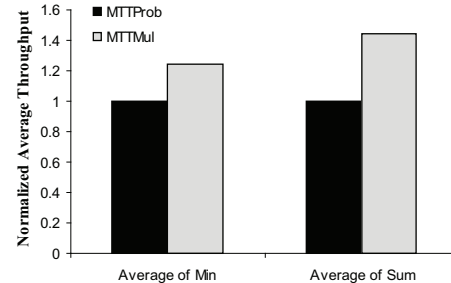


Figure 7: Normalized averages of the minimum and sum of each flow pair's throughputs for 20 flow pairs running on the MiNT- m testbed under MTTMul and MTTProb schemes. MTTMul improves minimum by 24% and the sum by 44%




Figure 8: Throughputs of 15 simultaneous UDP flows on a 5x8 simulated network under MTTMul and MTTProb schemes. MTTMul reduces to MTTProb for some flows where a single good path exists. When MTTMul chose two-path routes, it outperforms MTTProb in all cases except for Flow 6 and 11.

UDP and TCP-HH. In all cases, MTTProb outperforms MTTDelay and MTTBW. This is more so because the CDF based scheme is indeed able to better approximate a path's capacity than a mean value, which is less reliable when there is substantial link quality fluctuation. For UDP and TCP-HH, MTTBW outperforms MTTDelay as the adverse effect of packet losses on connection throughput is largely removed.

3.3.4 Multi-path Metric

Compared with MTTProb, MTTMul further improves the robustness of a wireless routing decision by leveraging diversity in multiple paths. However, using multiple paths has its cost, in the form of increased self interference and adverse impact of packet reordering. If the best path between a node pair exhibits a consistently high quality, using multiple paths to route traffic between that node pair would only degrade the effective throughput. Therefore, in its current design, MTTMul chooses a multi-path route if and only if

Scheme	Average (Mbps)	Aggregate (Mbps)	Fairness Index
MTTMul	0.225	3.4	0.308
MTTProb	0.192	2.91	0.197

Table 5: Average throughput, aggregate throughput and fairness index of MTTMul and MTTProb for the 15 flows for the experiment in Figure 8.

such a route’s estimated capacity is higher than that of the single best path found using MTTProb.

To compare MTTMul with MTTProb, we randomly chose 40 node pairs in the 3x4 grid, set up a UDP flow between each node pair, and ran two flows at a time. In each run, the two simultaneous flows are routed either both using MTTMul or MTTProb, and we measured the minimum (MIN) and the sum (SUM) of these two flows’ throughputs. The average results for the measurements of 20 such runs are shown in Figure 7. MTTMul outperforms MTTProb in both MIN and SUM measurements. MTTMul does better in MIN throughput because using multiple paths provides an effective hedge against sudden decrease in link quality especially for poorly connected node pairs. MTTMul also improves the aggregate throughput of the two test flows as it utilizes more network resources simultaneously.

To provide further understanding of how multi-path routing helps, we performed a series of trace-driven simulations. The network topology used in the simulations was a 5x8 (40-node) grid. The quality distribution of links in the simulated network was 10% good links, 46% medium links and 44% poor links, which corresponded to the link characteristics observed in the 3x4 MiNT-m testbed topology. Two of the 40 nodes were designated as gateways, and flows were set up between non-gateway nodes to one of the two gateways. We simulated 15 simultaneous UDP flows on the network and measured their throughputs. The results are shown in Figure 8, organized according to the number of paths used in the final route selected by MTTMul. MTTMul’s single-path route choice is always the same as MTTProb’s choice. For this topology MTTMul never uses more than two paths as the increase in self interference prevents it from choosing more than two paths. When MTTMul chooses two-path routes, it outperforms MTTProb in all cases except for Flow 6 and 11. This is to be expected because the 15 flows were sharing the entire network resource and the performance improvement for some flows may come at the cost of degraded performance of others. However, because MTTMul utilizes the network resources more efficiently and is able to exploit path diversity the aggregate throughput and the fairness index computed as $(\sum_i X_i)^2/n \times \sum_i X_i^2$, where X_i is the throughput of i^{th} flow, of these 15 flows are all higher under MTTMul than under MTTProb, as shown in Table 5.

4. RELATED WORK

There is a large body of literature on channel quality aware routing protocols. Due to limitations in space, we restrict ourselves to those that are most relevant in terms of different protocol mechanisms incorporated into CARP.

4.1 Wireless Routing Metrics

Since wireless link quality is affected by many factors, empirical techniques are best suited to estimate it. The ETT link metric was proposed for the WCETT routing metric [10] for the MCL routing protocol, which was designed for a multi channel network. WCETT of a n hop path is:

$$WCETT = (1 - \beta) \times \sum_{i=1}^n ETT_i + \beta \times \max_{1 \leq j \leq k} X_j. \quad (13)$$

$$X_j = \sum_{\text{Hop } i \text{ is on channel } j} ETT_i \quad 1 \leq j \leq k \quad (14)$$

Here the X_j component representing channel diversity, for a system with k orthogonal channels and β is a parameter that controls the weight of each term. For a single channel network WCETT metric reduces to $ETTDelay$ ($\beta = 0$). This however does not accurately account for contention from other flows in the neighbourhood and the suggested packet-pair technique to estimate link bandwidth, requires the use of explicit probe packets. Our results demonstrate that the MTT metric not only estimates the effective link and path bandwidth more accurately but also uses on-going transmissions to measure at a fine time scale. Note that MTT would work just as well even in a multichannel network since it would only account for those links which interfere in the same channel.

Several recent works [2, 23, 22] incorporate channel load or contention in the routing metric. The Medium Time Metric (MTM) [2] proposed by Awerbuch et al, uses the estimated end to end transmission time over a link. Estimated TT assumes that packets have fixed protocol overhead and channel access delay (back-off) is not included. The PARMA metric [23] uses passive channel sensing to measure the expected access delay of each packet and channel utilization to account for contention from other nodes. However passive sensing with current wireless technology requires the use of an additional wireless card. The ICTT metric by Zhai et al.[22] is proven to be close to an optimal routing metric and uses the transmission time along the maximal clique of the path instead of the transmission time along the most limiting link (MTTBW). However, in real-world networks ICTT is hard to measure as cliques change over time. MTTBW has the limitation that our estimation of I_l^p is an approximation for actual interference, However in our evaluations we see that this performs quite well.

One recent work is a channel quality measurement framework called EAR [12]. This work looks at best techniques to measure channel quality, and shows that broadcast based probes to measure channel quality do not accurately reflect the behaviour of unicast traffic over the same links. Also, broadcast probes fail to account for directionality [5] of wireless link quality. Hence it is better to use a directional metric and passive measurements when possible.

4.2 Diversity and Multipath Routing

Multi-path routing has been used in several mobile wireless network protocols [14, 15] to increase reliability of data transfer in mobile ad-hoc networks. Multi-path routing can also be used in wireless networks to exploit channel diversity. One approach to maximizing throughput is to leverage the broadcast nature of wireless transmissions. Specifically, when an intermediate receiver fails to receive a packet, one can leverage the fact that a neighbor of the intermediate receiver has successfully received the packet and could potentially forward it to the destination, thereby reducing retransmissions. ExOR is an opportunistic routing mechanism that makes use of the above intuition in making delayed forwarding decisions at the Routing layer[4]. Each relay node forwards packets only if it fails to overhear a higher priority node forwarding the same. However, it may

not always be possible for neighbours of a node to overhear each other, and hence multiple copies of the same packet could be propagated in the network. ROMER [21], another routing protocol based on this framework leverages transient variations to select the highest throughput path using a credit based forwarding scheme. ROMER forms an opportunistic, forwarding mesh centered around the long-term stable, minimum-cost path which opportunistically expands or shrinks at runtime to exploit the highest-quality links.

Both ExOR and ROMER exploit the redundancy due to the broadcast nature of the medium, but do not leverage the multirate option at the physical layer. Additionally, these schemes are not easily extensible to multi-channel networks. Multi-channel networks which provide much higher throughput [18], are usually designed such that neighbouring nodes are on different channels to reduce interference. In such a scenario the use of broadcast based multi-user diversity is limited. However, multi-path routing can still be exploited in multi-channel networks to increase throughput while using auto-rate mechanisms of modern cards.

5. CONCLUSION

Instead of comparing specific routing protocols, we take the approach of first defining the essential dimensions in the design space of WMN routing protocols, and then comparing the design alternatives incorporated into specific routing protocols in each dimension. Some important lessons have been learned and they are

- MTT is a more accurate wireless link capacity metric because it incorporates all performance factors including contention delay.
- Taking the bottleneck link's residual capacity is a more effective way to estimate a wireless path's capacity than summing up the delays of the path's constituent links.
- Using a probability distribution to model a wireless path's capacity is more robust in the presence of temporal fluctuation of wireless link quality, than using a single scalar value.
- Multi-path routing can improve the throughput robustness of wireless flows only if it is used when it is beneficial to do so.

6. REFERENCES

- [1] D. Aguayo, J. Bicket, S. Biswas, G. Judd, and R. Morris. Link-level measurements from an 802.11b mesh network. In *Applications, Technologies, Architectures, and Protocols for Computer Communication*, pages 121–132, 2004.
- [2] B. Awerbuch, D. Holmer, and H. Rubens. The medium time metric: High throughput rate selection in multi-rate ad hoc wireless networks. In *Kluwer Mobile Networks and Applications (MONET) Journal Special Issue on Internet Wireless Access: 802.11 and beyond*, pages 253–256, April 2006.
- [3] J. Bicket, D. Aguayo, S. Biswas, and R. Morris. Architecture and evaluation of an unplanned 802.11b mesh network. In *11th annual international conference on Mobile computing and networking (MOBICOM)*, pages 31–42, September 2005.
- [4] S. Biswas and R. Morris. Opportunistic routing in multi-hop wireless networks. In *ACM SIGCOMM Computer Communication Review*, pages 69–74, 2004.
- [5] A. Cerpa, J. L. Wong, M. Potkonjak, and D. Estrin. Temporal properties of low power wireless links: modeling and implications on multi-hop routing. In *6th ACM Mobihoc*, pages 414–425, 2005.
- [6] D. S. J. D. Couto, D. Aguayo, J. Bicket, and R. Morris. A high-throughput path metric for multi-hop wireless routing. In *9th ACM International Conference on Mobile Computing and Networking (MOBICOM)*, September 2003.
- [7] P. De, A. Raniwala, R. Krishnan, K. Tatavarthi, J. Modi, N. Syed, S. Sharma, and T. Chiueh. Mint-m: an autonomous mobile wireless experimentation platform. In *Mobisys*, pages 124–137, 2006.
- [8] P. De, A. Raniwala, S. Sharma, and T. Chiueh. Mint: A miniaturized network testbed for mobile wireless research. In *IEEE Infocom*, pages 2731–2742, March 2005.
- [9] R. Draves, J. Padhye, and B. Zill. Comparison of routing metrics for static multi-hop wireless networks. In *ACM SIGCOMM*, September 2004.
- [10] R. Draves, J. Padhye, and B. Zill. Routing in multi-radio, multi-hop wireless mesh networks. In *10th annual international conference on Mobile computing and networking*, pages 114–128, September 2004.
- [11] D. B. Johnson and D. A. Maltz. Dynamic source routing in ad-hoc wireless networks. In *Mobile Computing*. Kluwer Academic Publishers, 1996.
- [12] K.-H. Kim and K. G. Shin. On accurate measurement of link quality in multi-hop wireless mesh networks. In *12th annual international conference on Mobile computing and networking (MOBICOM)*, pages 38–49, 2006.
- [13] R. Krishnan, A. Raniwala, and T. Chiueh. Design of a channel characteristics-aware routing protocol. In *Proc. of IEEE Infocom Mini-Conference*, April 2008.
- [14] S.-J. Lee and M. Gerla. Split multipath routing with maximally disjoint paths in ad hoc networks. In *IEEE International Conference on Communication (ICC)*, pages 3201–3205, April 2001.
- [15] R. Leung, J. Liu, E. Poon, A.-L. C. Chan, and B. Li. Mp-dsr: A qos-aware multi-path dynamic source routing protocol for wireless ad-hoc networks. In *26th Annual IEEE International Conference on Local Computer Networks (LCN)*, page 132, April 2001.
- [16] C. E. Perkins and E. M. Royer. Ad hoc on-demand distance vector routing. In *Proc. of 2nd IEEE Workshop on Mobile Computing Systems and Applications*, pages 90–100, November 1999.
- [17] A. Raniwala, P. De, S. Sharma, R. Krishnan, and T. Chiueh. End-to-end flow fairness over ieee 802.11-based wireless mesh networks. In *IEEE Infocom*, pages 2361–2365, May 2007.
- [18] A. Raniwala, K. Gopalan, and T. Chiueh. Centralized channel assignment and routing algorithms for multi-channel wireless mesh networks. In *ACM SIGMOBILE Mobile Computing and Communications Review (MC2R)*, pages 50–65, April 2004.
- [19] A. Raniwala, S. Sharma, P. De, R. Krishnan, and T. Chiueh. Evaluation of a stateful transport protocol for multi-channel wireless mesh networks. In *IEEE Workshop on Quality of Service (IWQoS)*, pages 74–82, June 2007.
- [20] T. S. Rappaport. *Wireless Communications: Principles and Practice*. IEEE Press, Piscataway, NJ, 1996.
- [21] Y. Yuan, H. Yang, S. Wong, S. Lu, and W. Arbaugh. Romer: Resilient opportunistic mesh routing for wireless mesh networks. In *IEEE Workshop on Wireless Mesh Networks (WiMesh)*, September 2005.
- [22] H. Zhai and Y. Fang. Impact of routing metrics on path capacity in multirate and multi-hop wireless ad hoc networks. In *International Conference on Network Protocols (ICNP)*, November 2006.
- [23] S. Zhao, Z. Wu, A. Acharya, and D. Raychaudhuri. Parma: a phy/mac aware routing metric for ad-hoc wireless networks with multi-rate radios. In *Sixth IEEE International Symposium on a World of Wireless Mobile and Multimedia Networks, (WoWMoM)*, pages 286–292, June 2005.