The Medium Time Metric: High Throughput Route Selection in Multi-rate Ad Hoc Wireless Networks*

Technical Report
October 2004

Baruch Awerbuch, David Holmer, Herbert Rubens
{baruch, dholmer, herb}@cs.jhu.edu
Johns Hopkins University,
Department of Computer Science,
Baltimore, Maryland 21218 USA

ABSTRACT

Modern wireless devices, such as those that implement the 802.11abg standards, utilize multiple transmission rates in order to accommodate a wide range of channel conditions. The use of multiple rates presents a significantly more complex challenge to ad hoc routing protocols than the traditional single rate model. The hop count routing metric, which is traditionally used in single rate networks, is sub-optimal in multi-rate networks as it tends to select short paths composed of maximum length links. In a multi-rate network, these long distance links operate at the slowest available rate, thus achieving low effective throughput and reduced reliability due to the low signal levels. In this work we explore the lower level medium access control and physical phenomena that affect routing decisions in multi-rate ad hoc networks. We provide simulation results which illustrate the impact of these phenomena on effective throughput and show how the traditional minimum hop routing strategy is inappropriate for multi-rate networks. As an alternative, we present the *Medium Time Metric* (MTM) which avoids using the long range links often selected by shortest path routing in favor of shorter, higher throughput, more reliable links. Our experimental results with 802.11g radios show that the Medium Time Metric achieves significantly higher throughput then alternative metrics. We observed up to 17 times more end-to-end TCP throughput than when the Min Hop or ETX metrics were used.

Keywords

multi-rate, ad hoc, wireless, routing, routing metric, cross layer interaction.

1. INTRODUCTION

Ad hoc wireless networks are self-organizing multi-hop wireless networks where all nodes take part in the process of forwarding packets. One of the current trends in wireless communication is to enable devices to operate using multiple transmission rates. Many existing wireless networking standards include this multi-rate capability including IEEE 802.11abg[1]. The reason for this multi-rate operation stems directly from some of the fundamental properties of wireless communication. Due to the physical properties of communication channels, there is a direct relationship between the rate of communication and the quality of the channel required to support that communication reliably. Since distance is one of the primary factors that determines wireless channel quality, there is an inherent trade-off between high transmission rate and effective transmission range.

This range speed trade-off is what has driven the addition of multi-rate capability to wireless devices. Consumer demands for wireless devices always include both higher speed and longer range. Unfortunately a single rate represents a single trade-off point between these two conflicting goals. Since multi-rate devices support several rates, they provide a wide variety of trade-offs available for use. This gives them a great deal of flexibility to meet the demands of consumers. This added flexibility is the primary driving force behind the adoption of multi-rate capability. It is also reasonable to assume that this type of capability will also be present in future wireless networking standards.

While multi-rate devices provide increased flexibility, they cannot change the inherent trade-off between speed and range. Both high speed and long range cannot be achieved simultaneously. Long range communication still must occur at low rates, and high-rate communication must occur at short range. This multi-rate capability merely provides a number of different trade-off points. Multi-rate devices must have protocols that select the appropriate rate for a given situation.

In infrastructure based networks, all communication takes place between nodes and access points. In this case, an additional protocol required to support multi-rate is necessary only at the medium access control (MAC) layer. Single rate nodes already have the ability to select the best access point based on the received signal strength. Thus the only additional task necessary is that of selecting the actual rate used to communicate. Since the distance between the user and the access point is dictated by the physical geometry of the network, the rate selection task must react to the existing channel conditions. In other words, the only option available to a wireless device is to select the fastest modulation scheme that works reliably.

However, this is not the case in ad hoc multi-hop wireless networks. In these networks, the routing protocol must select from the set of available links to form a path between the source and the destination. While in single-rate networks all links are equivalent, in multi-rate networks each available link may operate at a different rate. Thus the routing protocol is presented with a much more complex problem. Which set of trade-offs does it choose? Long distance links can cover the distance to the destination in few hops, but then the links would be forced to operate at a low speed. Short links can operate at high rates, but more hops are required to reach the destination. In addition, the path selected by the routing protocol will not only affect the packets moving along that path, but will affect the level of congestion at every node within the interference range of the path as well.

Our Contribution. We provide an analysis of the lower level medium access control and physical phenomena that affect routing decisions in multi-rate ad hoc wireless networks. Based on this analysis we derive a general theoretical model of the attainable throughput in multi-rate ad hoc wireless networks. The traditional technique used by most existing ad hoc routing protocols is to select minimum hop paths. These paths tend to contain long range links that have low effective throughput and reduced reliability. We present the *Medium Time Metric* (MTM) [2] that selects higher throughput paths and tends to avoid long unreliable links. The MTM minimizes the total medium time consumed sending packets from a source to a destination. This results in an increase in total network throughput.

The rest of the paper is organized as follows. Section 2 summarizes related work. We further define our network model and assumptions in Section 3. In order to fully understand the effects of the physical and MAC layers on network throughput, we present a detailed analysis in Section 4. We examine existing route selection techniques in Section 5. The Medium Time Metric is presented in Section 6. In Section 7 we present a theoretical model of throughput in multi-rate networks and derive an optimal route selection heuristic. We examine the effects of node density in Section 8 and provide real world experimental results in Section 9. We then conclude in Section 10.

2. RELATED WORK

Ad Hoc Routing Protocols. A large number of routing protocols have been proposed by the ad hoc wireless networking community. Typically these have adopted one of two major strategies: on-demand such as in AODV [3] and DSR [4], and proactive such as in DSDV [5] and OLSR [6]. The vast majority of these protocols where originally designed for single-rate networks, and thus have used a shortest path algorithm with a hop count metric (min hop) to select paths. While min hop is an excellent criteria in single-rate networks where all links are equivalent, it does not accurately capture the trade-offs present in the more complicated multi-rate networks.

Signal Stability Based Ad Hoc Routing Protocols. In [7] the authors show that the minimum hop path generally contains links which exhibit low reliability. In [8] and [9] the authors present various routing protocols which are based on signal stability and link reliability rather then just shortest path in order to provide increased path reliability. In our work, signal information is used not only to increase path reliability, but also to increase network throughput.

Routing Metrics. De Couto et. al. present the Expected Transmission Count Metric (ETX)[10] that selects paths which minimize the expected number of transmissions required to deliver a packet from the source to the destination. The authors demonstrate through measurements on an experimental test-bed that links in an ad hoc network experience vastly different loss rates and that these loss rates fluctuate over time even in a static network [11]. Since the 802.11 MAC protocol retransmits lost packets, routing across lossy links significantly increases medium consumption and reduces throughput. The ETX metric attempts to measure link reliability and to select paths which minimize the expected total number of transmissions.

In [12], Draves et. al. provide a performance comparison of four proposed multi-hop routing metrics. Specifically: Hop Count, Per-hop Round Trip Time[13], Per-hop Packet Pair Delay[14], and Expected Transmission Count[10]. The evaluation was performed on a 23 node wireless test-bed using 802.11a wireless interfaces. Their results indicate that ETX outperforms Hop Count in a static network, but that Hop Count outperforms ETX under mobility.

Draves et. al. present Weighted Cumulative Expected Transmission Time (WCETT)[15], a routing metric for routing in multi-radio multi-hop static wireless networks. The authors refer to the combination of the Medium Time Metric (MTM)[2] with the Expected Transmission Count Metric (ETX)[10] as the Expected Transmission Time Metric (ETT). They propose an additional weighted factor β, which promotes channel diversity, and refer to the total combination as WCETT. The authors indicate that MTM+ETX provides a 16-55% increase in throughput over ETX alone and a 38.6% increase over minimum hop routing in a single radio environment. Their results in a multi-radio environment indicate that MTM+ETX achieves a median throughput increase of approximately 80% over ETX alone. An additional 10% throughput gain over MTM+ETX was achieved using the authors proposed WCETT channel diversity strategy.

A more detailed analysis of existing route selection techniques is provided in Section 5 and experimental results are available in Section 9.

Rate Selection. The method of rate selection in multi-rate capable networks has been left unspecified by the 802.11 standards. As a result, several auto rate protocols have been proposed. The most commonly used protocol is Auto Rate Fallback (ARF). ARF was originally developed for Lucent's WaveLAN II devices [16], and was later enhanced for 802.11b devices [17]. ARF operates using the link level ACK frames specified by the 802.11 standard. Each node increases the rate it is using to communicate with its neighbor after a number of consecutively received acks, and decreases the rate after a number of consecutively missed acks. The advantage of this technique is that it is easy to implement because it is purely sender based, requires no modifications to the 802.11 standard.

As an alternative, the Receiver Based Auto Rate (RBAR) protocol was presented in [18]. RBAR allows the receiving node to select the rate. This is accomplished by using the SNR of the RTS packet to choose the most appropriate rate and communicating that rate to the sender using the CTS packet. This allows much faster adaptation to the changing channel conditions than ARF, but requires some modifications to the 802.11 standard.

The Opportunistic Auto Rate (OAR) protocol, which is presented in [19], operates using the same receiver based approach, but allows high-rate multi-packet bursts to take advantage of the coherence times of good channel conditions. These bursts also dramatically reduce the overhead at high rates by amortizing the cost of the contention period and RTS CTS frames over several packets. By picking appropriate sized bursts, OAR also changes the fairness characteristic from each node sending an equal number of packets to each node getting an equal allocation of medium time. This produces a dramatic increase in overall throughput when links of multiple rates operate together in the same space. OAR also requires modifications to the 802.11 standard.

3. NETWORK MODEL

The multi-rate network model presented in this paper is based on the 802.11b standard [20]. The topics discussed here apply to other multi-rate standards, but all examples, ranges, and rates shown in this work are based on 802.11b.

Throughout the remainder of the paper we present the results of a number of NS2 [21] simulations. In order to simulate multi-rate 802.11b, we started with the ns-2.1b7a code base and the multi-rate extensions available from the Rice Networks Group [22] that contain implementations of the RBAR and OAR protocols. The 802.11 MAC and physical wireless parameters were further modified to match the published specifications of a Lucent ORiNOCO PC Card [23], a commonly used 802.11b wireless adapter (see Table 1). Since the carrier sense (CS) threshold specification is not published, we provide an estimate of 2.24 times the 1.0 Mbps receive range. This estimate is consistent with both the NS2 default carrier sense range, and the real world experimental results published in [24].

$$PathLoss(d) = \left(\frac{4\pi Hz}{c}\right)^2 d^4 \tag{1}$$

Table 2 shows the ranges resulting from these simulation parameters. The path-loss model in Equation 1 uses an exponent of 4, and was selected since it is representative of an indoor environment[25]. The ranges presented in this work are different then those presented in our previously work. The current ranges more accurately represent real worlds range, and are exactly proportional to the previously reported distances. The results presented here should be valid for any set of ranges with similar proportions regardless of magnitude.

4. MULTI-RATE THROUGHPUT

The total network throughput attainable in multi-rate ad hoc wireless networks is a result of the combined behavior of the medium access control protocol, routing protocol, and physical properties of a wireless network. In order to provide an understanding of how this combined behavior affects network throughput, we examine several different phenomena.

Table 1: 802.11b Wireless Card Parameters

Parameter	Value
Frequency	2.437 GHz
Transmit Power	15 dBm
11.0 Mbps Receive Threshold	-82 dBm
5.5 Mbps Receive Threshold	-87 dBm
2.0 Mbps Receive Threshold	-91 dBm
1.0 Mbps Receive Threshold	-94 dBm
Carrier Sense Threshold	-108 dBm

Table 2: 802.11b Ranges

Rate (Mbps)	Maximum Range
11.0	26.3 m
5.5	35.1 m
2.0	44.2 m
1.0	52.5 m
CS	117.7 m

4.1 Medium Access Control

Ad hoc wireless networks by nature use a broadcast medium. This means that any transmission made by a node simultaneously propagates to all other nodes in range. The downside of this property is that even if a node is sending packets to only one of its neighbors, those packets affect every other node in range. Furthermore, if two nodes transmit simultaneously, both transmissions will overlap and become garbled on the medium causing a receiver to be unable to successfully receive either packet. As a result, only a single transmission can occur at a time within range of the intended receiver.

The Medium Access Control (MAC) protocol is responsible for providing channel access arbitration and ensuring that nodes defer sending to avoid interfering with a transmission in progress. The 802.11 MAC protocol uses two mechanisms for deferral. The first mechanism used is carrier sensing, which means that the node listens to the medium in order to detect when another transmission is in progress. If it hears a transmission it defers until the medium is idle. Only nodes that are within carrier sense range of a sender will be able to successfully use this method to avoid collisions. The second mechanism is referred to as virtual carrier sense, and it is provided by a control frame exchange. A Request To Send (RTS) control frame is transmitted by the sender when it has a data packet to deliver. If the receiver is not already deferring, it responds with a Clear To Send (CTS) control frame. Any node that overhears an RTS or CTS is notified of the packet transmission, and will then defer for the duration of the transmission. This additional mechanism is particularly useful in cases where nodes near the receiver cannot carrier sense the transmission because of obstacles or other propagation effects. Figure 2 illustrates the ranges of these two mechanisms according to the specified communication model.

In addition to providing medium reservation, the RTS and CTS frames also serve other purposes. The first is fast collision resolution which is necessary because wireless devices are unable to use collision detection. The second is that the

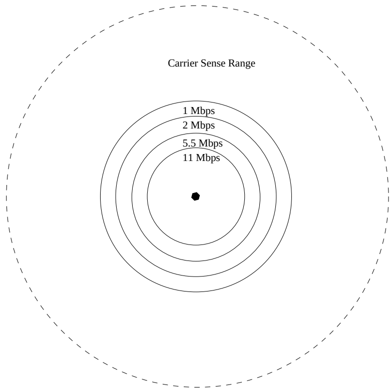

Figure 1: 802.11b Ranges

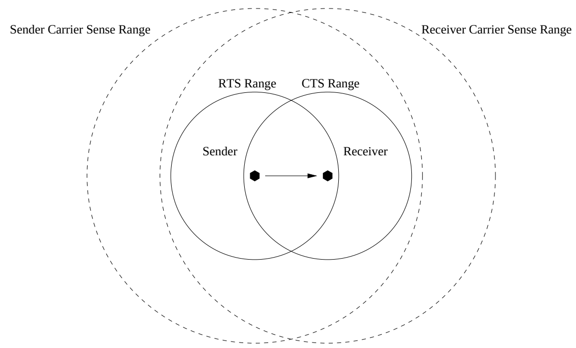

Figure 2: 802.11 Communication

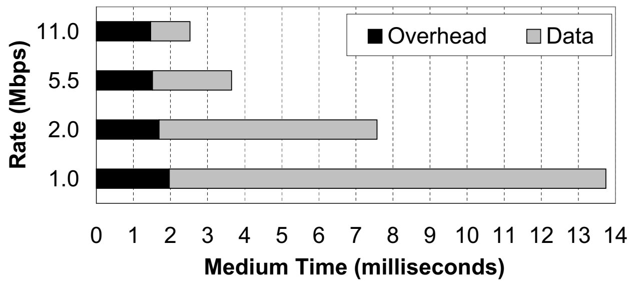

Figure 3: Medium Times for 802.11b Transmissions

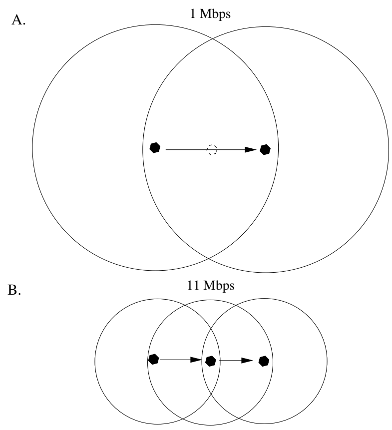

Figure 4: Path Selection Options

RBAR and OAR rate selection protocols use the RTS frame to provide a direct measurement of the current channel quality. The receiver can then select the most appropriate rate and notify the sender using the CTS frame. Since the receiver is able to select the rate every time it receives an RTS frame, it is able to respond quickly to variations in channel conditions.

The MAC protocol is responsible for providing channel access, which incurs a significant amount of overhead. In 802.11 this overhead is composed of three primary components: time spent transmitting control frames, random backoff time during contention resolution, and time wasted as a result of collisions.

Collision detection, which is used in Ethernet networks is impossible in wireless networks. In an effort to reduce the total overhead, 802.11 spends a significant amount of medium time sending control frames that are designed to help avoid costly data packet collisions. As a result, medium access control is more expensive in the wireless environment than in the wired environment.

The result of this MAC overhead is that the effective throughput is less than the link rate. Table 2 shows the results of a simple NS2 experiment where 1472 byte payload UDP packets were flooded across a single link using the full RTS-CTS-DATA-ACK exchange. The time spent for data and overhead in 802.11b are shown in Figure 3. The 802.11 MAC overhead is significant, particularly for the higher rate links. The effective throughput of an 11 Mbps link is less than half the link rate. Only the contents of the DATA and ACK frames are transmitted at the selected link rate, the rest of the exchange occurs at the 1 Mbps base rate. As a result, the MAC overhead is almost constant per packet. Therefore, the effective link rate is determined by the amount of time spent transmitting the data contents of each packet. We see a greater reduction in effective throughput for faster links because the time necessary to send a packet is inversely proportional to the rate of the link. In other words, because the data transmission time is small for fast links, the proportion of time consumed by the fixed overhead is large.

When considering the total throughput in the wireless network, it is important to consider the number of non-interfering transmissions that can simultaneously exist as well as the rate at which each transmission is occurring. Unfortunately, the number of simultaneous transmissions is determined by the physical network topology and the transmission power level. The greater the geographic size of the network the greater the number of possible simultaneous transmissions. A protocol cannot control the physical configuration of nodes in the network, but it can control the rate at which the nodes transmit data.

Given a network where three simultaneous transmissions can occur, if these transmissions are sent at 1 Mbps, which is the lowest 802.11b transmission rate, a maximum of 3 Mbps of total network throughput could be obtained. Consider the same network, but with transmissions occurring at 11 Mbps. This would result in a total network throughput of 33 Mbps, which is significantly greater.

4.2 Hops vs. Throughput Trade-Off

One approach to increasing throughput would be to configure all the nodes in the network to operate only at the highest transmission speed. This would ensure that the network would always operate at the maximum combined simultaneous rate. This approach may run into problems because of the inherent trade-off between the transmission rate and effective transmission range (see Figure 1).

In multi-hop ad hoc networks, packets must frequently traverse several hops to travel from the source to the destination. By using slow links that have high effective range, the distance between the source and destination can be covered using a small number of hops. If we avoid using all but the fastest links, we reduce the effective range of every node. One major drawback of this approach is that we run the risk of disconnecting components of the network. Even if we do not disconnect the network, we increase the number of hops required to cover the distance from the source to the destination.

Consider the following example where the source and destination are barely within range of one another (see Figure 4-A). In this configuration the source can reach the destination in one transmission at the lowest rate. A single link is by definition the minimum hop path between the source and the destination since no other path can be shorter. While sending the packet directly to the destination would result in the least number of transmissions, the transmission would occur at the slowest possible speed, requiring all of the other nodes in this neighborhood to defer transmitting for the longest possible time. As we previously discussed, transmitting at this rate will limit the overall throughput attainable in the network.

Now consider the same situation except an additional node is located between the source and the destination (see Figure 4-B). The source and destination can still communicate directly through one low speed transmission, but now an additional option exists. The traditional minimum hop path algorithm would not consider this configuration any differently from the previous, since routing through the intermediary node would only increase the hop count. The speed of each of the two transmissions would be 11 Mbps as opposed to the single 1 Mbps transmission selected by the minimum hop approach. This would provide an effective bandwidth along the path of 2.38 Mbps by utilizing two 11 Mbps hops as opposed to 0.85 Mbps across the single 1 Mbps link. This represents almost a three fold increase in throughput (see Tables 2 and 3).

The previous example suggests that choosing routes that use high-rate links is strictly better then those that use low-rate links. While this is true in many individual situations (including the one above), there are other factors to consider. In the previous example, two 11 Mbps links were used to provide increased throughput over the single 1 Mbps link. Despite the fact that all of the links in the path operate at 11 Mbps, the throughput of the path is only a fraction of a single 11 Mbps link. This is because only a single transmission can occur at a time within the same area. For the packets to traverse the two 11 Mbps hops, the source would have to alternate with the forwarding node. In other words, the nodes need to take turns transmitting. This coordination is handled by the medium access control layer.

In this simple example, the two 11 Mbps hops are strictly better than the single 1 Mbps hop, but this might not be the case if the choice is between ten 11 Mbps hops and a single 1 Mbps hop. There are several reasons why this is true. When packets are sent along a path in a multi-hop network, the adjacent transmissions are competing for access to the medium. By sending across many hops, the throughput along the path becomes a fraction of the capacity of the links. In Figure 5 nodes 1 and 8 are communicating along a path. The diagram shows the nodes that are affected by the transmission of node 4 while it is forwarding the packet on to node 5 along the path. In this example all eight nodes are being affected by the single transmission that is taking place and they all must defer from sending until the transmission completes.

Table 3: Two Hop Path Throughput

Link Rate (Mbps)		Path Throughput
1^{st} Hop	2^{nd} Hop	(Mbps)
11.0	11.0	2.38
11.0	5.5	1.86
11.0	2.0	1.15
5.5	5.5	1.59
5.5	2.0	1.04
2.0	2.0	0.77

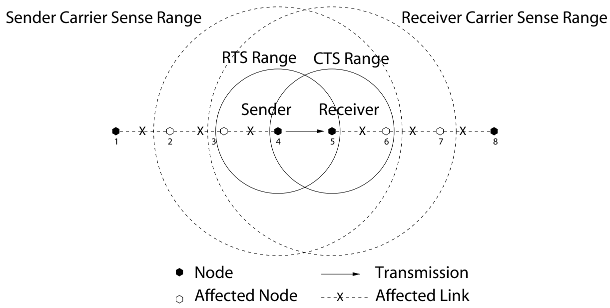

Figure 5: Effect of Transmission on Other Nodes

In this example, nodes 2 through 6 are all in carrier sense range of node 4, which is transmitting. These nodes all defer until the transmission completes. Node 7 on the other hand is in carrier sense range of the receiver but not the sender. Node 7 can carrier sense the receiver's CTS packet, but will not be able to carrier sense the actual transmission. This will cause node 7 to defer for an extended inter-frame spacing, which may not be long enough for the transmission from node 4 to 5 to complete. If node 7 begins transmitting it could potentially cause a collision. This example shows that the 802.11 MAC protocol has not solved the hidden terminal problem [26]. Another interesting aspect of this example is the effect of the transmission on nodes 1 and 8. Both of these nodes are out of the carrier sense ranges of both the sender and the receiver (nodes 4 and 5 respectively). As a result they appear to be unaffected by the transmission that is taking place. While it is true that these nodes could communicate with any other node outside of the current transmission neighborhood, in this particular example they are attempting to communicate along the path between node 1 and 8. Since nodes 2 and and 7 are currently deferring as a result of the transmission, any RTS initiated by nodes 1 or 8 would receive no reply. As a result nodes 1 and 8 will also need to defer until the transmission from node 4 to 5

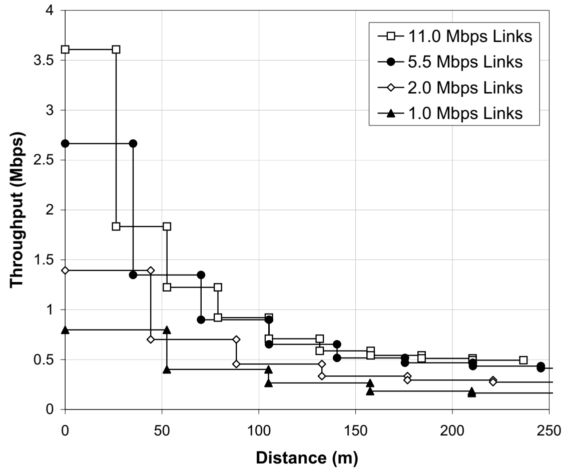

Figure 6: Throughput Loss Along a Path

completes. This example shows the broad impact that a single transmission has on nodes along the path as well as on other nodes in the immediate vicinity.

4.3 Quantitative Evaluation of Path Throughput Loss

An additional example shows a more quantitative evaluation of the throughput loss along a path. Figure 6 contains the results of a simulation that was conducted to explore the throughput loss of a single TCP connection along a path where each link operates at the same rate. Simulations were conducted for each of the four 802.11b link rates. The results show the throughput across the path vs. the distance (or length) of the path. As the length of the path increases the number of hops required to traverse the distance also increases. Since the throughput drops as the number of hops increases, the throughput drops in steps. The width of each step is equal to the effective transmission range at the given rate.

Since high-rate links have a shorter effective range, a greater number of hops is required to cover the same distance as a smaller number of lower rate hops. This is indicated in the graph since the high-rate throughput drops multiple times for each decrease in the low-rate throughput. There are a couple of interesting observations that are evident in this graph. The first observation is that the lines intersect. This means that at certain distances more throughput can be obtained using lower speed links then higher speed links. A specific example of this occurs at 30m. Notice the throughput obtained by the 5.5 Mbps path is greater than that of the 11 Mbps path. This occurs because the 11 Mbps path needed to traverse 2 hops at this distance, while the 5.5 Mbps path still consists of a single hop. This shows that traversing high speed links does not always achieve the highest throughput in all cases. Another interesting observation is that after approximately 175m the speeds seems to plateau. This is due to spatial reuse. As the path becomes longer, multiple transmissions can take place simultaneously along the path. This allows the throughput to reach a steady state, where additional distance does not cause any significant decrease in throughput. It is also important to notice that at this distance the throughput of the links increases as the link speed increases. This suggests that even though high link rate paths must traverse more links to reach the same distance, they still provide more throughput.

Table 4: Temporal Fairness Throughput Results

	Packet Fairness RBAR (Mbps)	Temporal Fairness OAR (Mbps)
11.0 Mbps Link	0.896	3.533
1.0 Mbps Link	0.713	0.450
Total	1.609	3.983

4.4 Temporal Fairness

In addition to low path throughput, there are other detrimental effects of sending packets at slow transmission speeds. The standard 802.11 MAC protocol attempts to provide fairness to individual senders on a per packet basis. This means that if there are two senders near each other and they are continuously trying to send packets, they should end up sending approximately the same number of packets. In multi-rate networks, there is no guarantee that these two senders are sending at the same rate. Since the MAC protocol is only attempting to be fair with regard to the number of transmissions, slow senders dominate the medium time. This effect has been demonstrated in [27] through both simulation and experimentation. One technique for dealing with this problem involves redefining the MAC fairness model. Temporal fairness would provide an equal share of medium time between senders independently of their transmission rate. There has already been work which explores this option.

The Opportunistic Auto Rate (OAR) protocol[19] provides temporal fairness with regard to medium time by allowing senders who send at a high-rate to send as many packets as required to equal the transmission time of a single packet at a low-rate. Basically, this results in every sender having an equal opportunity to transmit and for each sender to be able to transmit for the same amount of medium time. This is a dramatic improvement in efficiency over the existing 802.11 fairness model.

A simulation was run in NS2 to illustrate this effect. The simulation consisted of two nodes flooding packets to two different destinations. One sender was sending at 1 Mbps and the other was sending at 11 Mbps. All nodes in the simulation were within range of each other and were contending for access to the medium. The simulation was conducted with both the OAR and RBAR protocols and the average results are shown in Table 4.

As seen in the results, the OAR provides almost two and a half times the total throughput of RBAR. This indicates that temporal fairness is extremely important for achieving high throughput in ad hoc networks. The RBAR results, which are representative of the current 802.11 MAC, indicate that even if some of the routes in the network are operating at high link speeds, the total network throughput will still be low as a result of low speed links dominating network medium time. We conclude that in order to achieve high throughput, not only will the routing protocol need to be selecting high speed links, but the medium access control protocol will have to provide temporal fairness to ensure that low speed links do not gain an unfair share of the medium time.

5. ROUTE SELECTION METRICS

A number of route selection metrics have been proposed for path selection in multi-hop wireless networks. Most existing route selection metrics were designed for single rate networks. We examine existing metrics and explore their performance properties in multi-rate networks.

5.1 Minimum Hop Route Selection

Most existing ad hoc routing protocols have utilized hop count as their route selection criteria. This approach minimizes the total number of transmissions required to send a packet on the selected path. This metric is appropriate in single-rate wireless networks because every transmission consumes the same amount of resources. However, in multi-rate networks this technique has a tendency to pick paths with both low reliability and low effective throughput.

Throughput Loss. In multi-rate wireless networks, the selection of minimum hop paths typically results in paths where the links operate at low rates[7]. This is because the shortest path contains the fewest number of nodes between the source and destination. Fewer intermediate nodes corresponds to longer links in order to cover the same distance. Since distance is one of the primary factors that determines channel quality, the long links have low quality, and thus operate at low rates. So given the opportunity, in an effort to minimize the number of hops, shortest path selection protocols will pick paths composed of links close to their maximum range that must operate at the minimum rate.

Not only do the low link rates produce a low effective path throughput, but as a result of the shared wireless medium, this path selection degrades the performance of other flows in the network. This occurs due to the large amount of medium time required to transmit a packet at a slow link speed. All nodes within interference range of the transmission must defer while it takes place. Thus, slow transmissions reduce the overall network throughput by consuming a large amount of medium time.

Reliability Loss. Multi-rate wireless devices are inherently designed to deal with changes in channel quality due to mobility and interference. The devices provide multiple link speeds to accommodate fluctuations in link quality. In 802.11b, as two nodes move in opposite directions, the auto rate protocol will gracefully reduce their link speeds from 11 Mbps down to 1 Mbps before they are finally disconnected.

Minimum hop path route selection has a tendency to choose routes that utilize the lowest link speed, leaving the auto rate protocol no flexibility in dealing with channel quality fluctuations. As a result, routes are often established between nodes that are on the fringe of connectivity. The authors of [9] refer to this as the Grey Zone. This occurs when nodes are able to receive broadcast transmissions, but data/ack packets are unable to be successfully delivered. While routing broadcasts are typically extremely small in size, data packets typically occupy the full frame size, making them more susceptible to corruption at high bit error rates (BER). This tendency is even further exaggerated by the way 802.11 handles broadcast transmissions as opposed to unicast transmissions. While broadcasts are sent as a single frame, unicasts require a full RTS-CTS-DATA-ACK exchange for successful delivery, which is more likely to be disrupted by a low quality channel. The end result is that small broadcasts can often be delivered even when data communication is not possible.

5.2 Expected Transmission Count Metric

De Couto et. al. present the Expected Transmission Count Metric (ETX)[10] that selects paths which minimize the number of transmissions required to deliver a packet from the source to the destination. In a completely reliable network the ETX and Minimum Hop paths are exactly the same. However, the authors demonstrate through measurements on an experimental test-bed that links in an ad hoc network experience vastly different loss rates and that these loss rates fluctuate over time even in a static network [11]. Since the 802.11 MAC protocol retransmits lost packets, routing across lossy links significantly increases medium consumption and reduces throughput. The ETX metric attempts to measure link reliability and to select paths which minimize the expected total number of transmissions and retransmissions.

Throughput Loss. While it can be argued that the ETX metric is ideal in single rate networks (which it was designed for), multi-rate networks present an additional set of challenges. In multi-rate networks, the ETX metric is unable to determine the difference between an 11 Mbps and a 1 Mbps link which are experiencing the same loss rates. In addition, since high throughput paths are longer then minimum hop paths, the ETX metric will avoid them since they require more transmissions. This results in a loss of throughput in multi-rate networks since ETX is unable to properly take advantage of the higher throughput options which exist.

Reliability Loss. The ETX metric provides a considerable improvement in path reliability over minimum hop route selection in single rate networks. However in multi-rate networks, by minimizing the number of transmissions it has a tendency to select routes that utilize the lowest link speed. This leaves the auto rate protocol no flexibility in dealing with channel quality fluctuations, resulting in reduced path reliability.

Mobility. The ETX concept of selecting routes based on reliability is an excellent step forward over previous route selection techniques. However, ETX is currently implemented by actively probing and averaging reliability over a window of time (e.g. 10 seconds). Existing implementations have been shown to perform poorly under mobility. The authors of [12] suggest that ETX was unable to react fast enough to changes in link quality. It should be noted however that the ETX metric was designed to operate in static fixed wireless networks, not in mobile ad hoc networks.

6. MEDIUM TIME METRIC

We propose the *Medium Time Metric* (MTM) which is an additive metric designed to allow any shortest path routing protocol to select a high throughput path. The MTM assigns a weight to each link in the network that is proportional to the amount of medium time used by sending a packet on that link. Therefore the weight of any given path is the total medium time consumed when a packet traverses the whole path.

More formally we define $\tau(e, p)$ as the time required to transmit a packet p over edge e. $\tau(e, p)$ should take into account any sources of overhead such as contention, headers, and multiple frame exchanges. Given $\tau(e, p)$, the Medium Time Metric of a path π_{ij} for a packet p is defined as

$$MTM(\pi_{ij}, p) = \sum_{\forall e \in \pi_{ij}} \tau(e, p) \quad (2)$$

As a result, shortest path protocols that use the medium time metric find paths that minimize the total transmission time. The inverse of the medium time of a link is proportional to the "real" capacity of that link. Similarly, the inverse of the path MTM approximately equals the end-to-end path capacity. Therefore, a protocol using MTM simultaneously minimizes its usage of the shared medium and maximizes its end-to-end path capacity.

Medium Time Metric characteristics:

- By minimizing medium time consumption, path capacity is maximized.
- Minimizing medium time consumption, maximizes residual capacity available to other flows.
- Tracks path capacity as opposed to path utilization, thus is not prone to oscillation.
- Increases path elasticity under mobility.

6.1 Link Transmission Time

In order to compute the Medium Time Metric we must be able to estimate the time required to transmit a packet over a link $\tau(e, p)$. The Medium Time Metric originally considered only the effects of link rates [28], however more recent work by De Couto et. al. on the ETX metric shows that link reliability should also be taken into account [10]. Since wireless links are not completely reliable, a packet may need to be transmitted more then once, consuming additional medium time, in order to be successfully received.

We define $overhead(e)$ as the amortized average per packet overhead of a link including control frames, contention backoff, and fixed headers. $rate(e)$ represents the selected transmission rate, $size(p)$ represents the size of the data payload, and $reliability(e)$ is the fraction of packets which are successfully received. The link transmission time is

$$\tau(e, p) = \frac{overhead(e) + \frac{size(p)}{rate(e)}}{reliability(e)} \quad (3)$$

6.2 Estimating Link Transmission Time

In order to select a path which minimizes medium time consumption, techniques must be employed to estimate link transmission times. The link transmission time is composed of four discrete components: fixed overhead, packet size, link rate, and reliability estimation. The *link overhead* is calculated according to the specifications in the wireless standard and specifications of optional manufacturer provided features such as fast framing and packet bursting. It may include RTS, CTS, ACK, preamble, contention time, and any other sources of fixed overhead. This time will depend on both the type of wireless device and its configuration. A routing protocol should be able to query its wireless card's configuration parameters programmatically. The *packet size* should also be readily available to the routing protocol.

All multi-rate wireless devices provide an auto-rate mechanism for selecting a link rate for each of their neighbors. The technique utilized for auto-rate selection is unspecified by the 802.11 standards; the exact strategy is proprietary and varies between card manufacturers. Auto-rate selection and neighbor tracking is generally performed in the wireless card's firmware or device driver. In order for the auto-rate protocol to work efficiently, it must gather as much physical channel information as possible. This information might include: loss rate history, signal level, noise level, demodulation performance, or channel impulse response (as in MIMO systems). More advanced techniques will yield more accurate channel estimation. While this information is currently only used by the medium access control protocol, exporting this information to the routing protocol enables accurate estimation of both *link rate* and *link reliability*.

The most appropriate method for estimating link transmission times is to leverage information which is already being collected by the MAC and Physical layers. An alternate technique used in [10] and [15] is to perform active probing at the Network layer in order to measure loss rates and estimate link speeds. This approach is unable to take advantage of the more advanced channel quality estimators which are available at the lower layers. In addition, active probing techniques introduce additional network overhead proportional to the accuracy and rate at which they gather information. This makes them less suited for mobile environments. In this work, we strongly advocate inter-layer communication; particularly between the MAC and Network layers. In order to enable this, wireless radio manufacturers will need to provide a standard interface allowing higher layer protocols access to the neighbor state information which they are already maintaining[1]. This information is best maintained by the lower layers, and attempts to gather it at the network layer result in additional overhead and less accurate estimations.

6.3 Path Elasticity

One property of MTM paths is that they prefer high capacity links. Since the capacity of a link is directly related to the channel quality, high capacity links are able to absorb a channel quality reduction by lowering their rate. In contrast if the channel quality of a low capacity link is reduced, it will result in a link break. Channel quality reductions occur due to both mobility and environmental changes. For example, as two nodes move apart, the auto rate protocol gradually reduces the link speed. Nodes connected by a high-rate link must move a considerable distance before the link breaks. This allows routing protocols which utilize the MTM metric to select paths which are more elastic under mobility. This is particularly useful for protocols such as TCP or VoIP, since a path break results in a significant disruption to the protocol.

Simulations were conducted to evaluate the effects of path elasticity. In each simulation, 100 nodes were randomly placed in a 210m by 210m area[2]. The nodes move according to a random way-point mobility model with a maximum speed of 4 m/s. The node ranges are specified in Table 2. No routing protocol was used during these simulations. Instead, static all-to-all shortest path routes were computed at the beginning of each simulation using both minimum hop

9[1]The authors have already initiated a dialog with radio manufacturers and will continue to pursue this direction.

9[2]The simulations are approximately equivalent to a 1000m by 1000m area with nodes moving at a maximum speed of 20 m/s if NS2 default ranges are used.

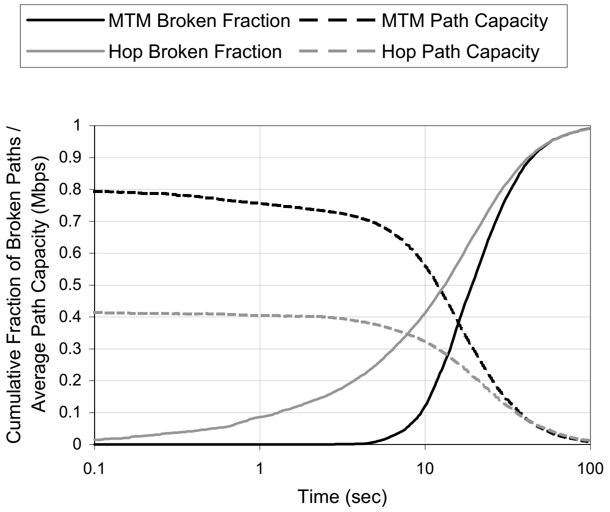

Figure 7: Path Elasticity Results

and MTM metrics. Over the course of the simulation, since the nodes are in motion, the paths eventually break. Figure 7 shows both the cumulative fraction of broken paths versus time and average path capacity versus time.

The elasticity is primarily shown by the difference between minimum hop and MTM in their cumulative fraction of broken paths over time. The results show that 10% of the minimum hop paths have broken after 1.3 seconds whereas 10% of the MTM paths have not broken until 9.5 seconds. If we look at how many paths have broken after 5 seconds we see that 25% of minimum hop paths have broken where as only 0.8% of MTM paths have broken. This clearly shows that MTM paths are significantly more elastic under mobility then minimum hop paths.

If we look at the average end-to-end path capacity versus time, we see that the average capacity begins at its maximum and goes to 0 as all of the paths are broken. The MTM paths begin the simulation with almost twice the capacity of the minimum hop paths. As the nodes begin moving, the MTM links begin to stretch and as distance increases their rates reduce. This is visually indicated by the downward slope of the MTM capacity line. After 5 seconds, only 0.8% of the MTM paths have broken, and the average path capacity has only decreased by 13%.

7. GENERAL MODEL AND OPTIMALITY ANALYSIS

There is some ambiguity in the literature regarding what constitutes an optimal solution for the routing problem in multi-hop wireless networks. One of the main reasons for this is the inherent difficulty in modeling the complex environment of wireless multi-hop networks. We provide a model that captures many of the effects present in such a network. Our model is similar to one presented in [29].

7.1 General Model of Attainable Throughput

In this work, we do not consider packet scheduling issues and consider a steady-state flow model. In this model, each network edge may be fractionally shared by several flows; however, the sum of shares cannot exceed 100%. Our model of the wireless network is defined by a *transmission graph* and *interference graph*.

The *transmission graph* is defined as $G(V, E, \rho)$. V is defined as the set of nodes in the network. A transmission edge $(u, v) \in E$ if node u is capable of transmitting to node v. ρ is a function that assigns a transmission rate to each transmission edge $\rho : E \rightarrow R^+$. $\rho(e) = \hat{\rho}$ where $\hat{\rho}$ is the maximum flow rate obtainable over edge e when no other traffic exists in the network. $\hat{\rho}$ should take into account any sources of overhead such as contention, headers, and multiple frame exchanges, and represents the "real" capacity of edge e. In this general definition, the transmission graph may be directed, and the transmission rate in the reverse direction of a bi-directional edge may be different than that in the forward direction. This is possible in real wireless networks because of different node configurations and asymmetric channel effects.

The *interference graph* is defined as $G(\tilde{V}, \tilde{E})$. We define the vertices of the interference graph to be the edges of the transmission graph, so $\tilde{V} = E$. An edge in the interference graph represents the interaction between packets transmitted on nearby transmission edges. $((a, b), (c, d)) \in \tilde{E}$ if $(a, b), (c, d) \in E$ and if a transmission on (a, b) interferes with a transmission on (c, d).

Given the interference graph, we can define the interference neighborhood of any given edge (u, v) as follows.

$$\chi(u, v) = \{(u, v)\} \cup ((x, y) : ((x, y), (u, v)) \in \tilde{E}) \quad (4)$$

Consider a set of i flows, where each flow ϕ_i originates from source s_i and is sinked by receiver r_i. Without loss of generality, we can represent each flow as a sum of path flows (indexed by j).

$$\phi_i = \sum_j \phi_{ij} \quad (5)$$

Each path flow ϕ_{ij} exists only on π_{ij}, where π_{ij} is a path from s_i to r_i in the transmission graph. In other words, $\phi_{ij}(x, y)$ equals the magnitude of the path flow $|\phi_{ij}|$ if the edge lies on its path, $(x, y) \in \pi_{ij}$, or zero otherwise. Thus we have effectively decomposed the general flow ϕ_i which may traverse multiple paths simultaneously into a set of flows ϕ_{ij} that each traverse only a single path, but sum to the original flow ϕ_i.

With this setup, we can now specify a flow constraint that captures the phenomena discussed above. For each edge (u, v) in the transmission graph, the sum of the fractional shares used by all flows in the interference neighborhood of (u, v) must be less than or equal to 100%. This is a more complicated version of the classic edge capacity flow constraint.

$$\sum_{(x,y) \in \chi(u,v)} \sum_{i,j} \left(\frac{\phi_{ij}(x, y)}{\rho(x, y)} \right) \leq 1 \quad (6)$$

In this general case, Linear Programming (LP) methods are required to achieve an optimal throughput solution. Opportunity-cost based approximations are possible in both the off-line case [30] (all connections are known ahead of time) and in the online case [31, 32]. Single path solutions are even harder to achieve as they require integer LP approaches.

7.2 Optimal Routing Assuming a Complete Interference Graph

Consider the special case of the general model where the interference graph is a clique (completely connected graph), i.e. each node can carrier sense each other node. In this special case, the constraint can be simplified since the interference neighborhood of any edge $\chi(u, v)$ is the same and consists of every edge in the transmission graph. In this case we wish to show the following theorem:

THEOREM 1. *In the case of a complete interference graph in the stated multi-rate ad hoc wireless network model, a routing protocol that chooses a single path that minimizes the sum of the transmission times optimally minimizes network resource consumption, and optimally maximizes total flow capacity.*

Given the complete interference condition, we can rewrite the general flow constraint.

$$\sum_{(x,y)\in E} \sum_{i,j} \left(\frac{\phi_{ij}(x,y)}{\rho(x,y)} \right) \leq 1 \tag{7}$$

We can reverse the order of summation.

$$\sum_{i,j} \sum_{(x,y)\in E} \left(\frac{\phi_{ij}(x,y)}{\rho(x,y)} \right) \leq 1 \tag{8}$$

We can also decompose $\phi_{ij}(x, y)$ by moving its magnitude out of the inner sum, and changing the inner sum to include only non-zero terms.

$$\sum_{i,j} \left(|\phi_{ij}| \cdot \sum_{(x,y)\in \pi_{ij}} \left(\frac{1}{\rho(x,y)} \right) \right) \leq 1 \tag{9}$$

Since $\rho(x, y)$ was defined as the real capacity of transmission edge (x, y), we can define the transmission time used by a unit of flow on this edge to be the inverse of this capacity.

$$\tau(x,y) = \frac{1}{\rho(x,y)} \tag{10}$$

Thus the final constraint equation becomes

$$\sum_{i,j} \left(|\phi_{ij}| \cdot \sum_{(x,y)\in \pi_{ij}} \Big(\tau(x,y) \Big) \right) \leq 1 \tag{11}$$

In other words, the flow over each sub path consumes a certain fraction of the capacity. The sum of these fractions must be less than one. The fraction consumed by each sub path is equal to the amount of flow on that path times the sum of the transmission times along that path. The magnitude of flow on a sub path, $|\phi_{ij}|$, will be maximized when the sum of the transmission times along that path, $\sum_{(x,y)\in\pi_{ij}} \tau(x, y)$, is minimized. Therefore, a routing protocol that selects paths that minimize the sum of the transmission times maximizes the flow along those paths. Also, it is only necessary for each flow to have a single sub path that minimizes the sum of the transmission times, because any other sub paths will be at best equivalent to the minimum, and thus offer no additional flow capacity. Even if a flow does not use its maximum available capacity, minimizing the path transmission time minimizes the flow's consumption of the common network resource and allows other flows to increase. Thus we have shown Theorem 1 to be true.

7.3 Optimality Discussion

We have shown that the MTM is globally optimal when all of the links in the network interfere with each other. In the general case, modeling the interference graph of an arbitrary network may be quite difficult due to complex propagation effects caused by obstacles and reflections. In real networks the interference graph is primarily determined by the carrier sense range. The interference graph includes "edges" between each possible transmission edge, and all other transmission edges with an endpoint within carrier sense range of one of the transmission edge's endpoints. While the carrier sense range is not infinite, in 802.11b networks experimental results show that it is greater that twice the maximum transmission range[24]. This roughly means that a transmission will interfere with every node within a two hop neighborhood.

Multi-hop Access Point Model. The simplest example of a complete interference network would be a group of nodes associated with an access point. Since all of the nodes are in transmission range of the access point, they are all able to carrier sense each other. In this configuration only a single transmission can occur at a time. While multi-hop routing is generally not considered when discussing access point connectivity, recent work such as [33] and [34] present systems that provide multi-hop infrastructure access.

Access points are currently the most commonly used communication model for wireless networking. The fact that MTM is optimal for this application further motivates the metric as well as the significance of our theoretical result. In this work, we have demonstrated that in multi-rate networks routing across multiple high-rate hops can achieve higher throughput than across a single low-rate link. Our analytical results prove that the Medium Time Metric will deliver globally optimal network throughput in a multi-hop access point system model. The maximum throughput gains achieved by MTM will be the difference between the lowest transmission rate and the fastest possible MTM path achievable at that distance. Applying our results from Section 4, the MTM could gain up to 3 times the throughput in 802.11b networks and up to approximately 10 times the throughput in 802.11g networks[3]. These results strongly motivate the need for multi-hop routing even for traditional access point connectivity.

Pipelining Effect. In larger networks, we can no longer claim that the MTM is globally optimal because traffic patterns and congestion may shift the optimal routes. However, the MTM still exhibits excellent characteristics in these larger networks. This occurs because the sum of the medium times is an accurate predictor of total end-to-end path capacity, until the paths grow long enough to exhibit significant pipelining. This will not occur as long as one link in the path is in interference range of all other links in the path. In the current network model, this occurs with paths of up to seven maximum length hops (see Figure 5). Once the paths

9[3]Based on Atheros 802.11G specifications.

are long enough to exhibit significant pipelining, the MTM begins to underestimate their throughput potential.

The reader should note that this capacity estimation property places no restriction on the total network size and only restricts the length of the actual communication paths. This is an important observation since prior research has shown that ad hoc networks scale only if the traffic patterns remain local [35]. A local traffic pattern, such as when every node accesses the nearest Internet gateway, provides a natural path length limit allowing the MTM to operate accurately even in extremely large networks. Non-local traffic patterns do not scale even with a globally optimal routing protocol since the attainable throughput at the pipelining distance is extremely small. Thus communicating over a large number of hops is inefficient; a large quantity of medium time is consumed in exchange for a small quantity of end-to-end throughput.

8. NODE DENSITY SIMULATION

Routing protocols that use the medium time metric choose paths that minimize the total consumed medium time. Our results show that these paths yield significant throughput gains when compared with minimum hop paths. However, this assumes that a path exists that utilizes less medium time than the minimum hop path. This may not be the case. Whether a high throughput MTM path exists depends solely on the current network topology. If relay nodes are located in ideal locations, the metric will be able to extract the most performance. In general, the likelihood that nodes exist in these locations increases as the density of the network increases.

When the density of the network is low, the topology is sparsely connected. This yields few choices for routing protocols to select from. In this situation, MTM and min hop will tend to select the same path. Conversely, as the network density increases, the abundance of nodes creates a dense, heavily interconnected topology. Routing protocols are provided with a multitude of paths from which to choose. This large number of choices allows the natural tendencies of each metric to be fully expressed.

Table 5 shows the expected medium times, and corresponding weights, for each rate computed according to the 802.11b standard specifications. The times are calculated assuming a full RTS, CTS, DATA, ACK exchange. These computed times also include an estimate of the time spent backing off during contention. Once the link rates are known, then integer weights 5, 7, 14, and 25 are used for link rates 11, 5.5, 2, and 1 respectively.

We have constructed a simple experiment designed to illustrate the relationship between density and the performance of the MTM. A variable number of nodes are randomly placed along a straight line path of fixed length. A single UDP flow is setup between the source and destination, which are placed at opposite ends of the line. Figure 8 shows the relative throughput of the MTM and min hop routing protocols as the number of nodes and the line length are varied. The vertical axis shows the percent increase in achieved throughput over the min hop path when using the MTM. The horizontal axis shows the normalized density of the topology. We define the *normalized density* as the average number of nodes within the maximum transmission range of a given node.

The results show a clear relationship between node density and increased throughput. As expected, at low densities we see low increases as both the MTM and min hop metric select nearly the same path. As the density increases, we see the full potential of the MTM revealed. The MTM path yields greater than three times (+200%) the throughput of the min hop path with the higher densities and longer path lengths. Longer paths yield more increased throughput than shorter paths because the MTM path utilizes the extra medium time available in long paths (from spatial reuse) much more efficiently than the min hop path.

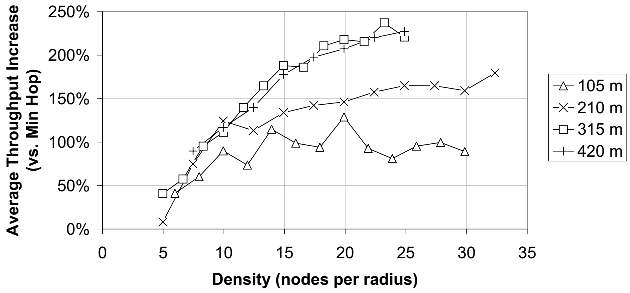

Figure 8: Average throughput increase of MTM along a randomized straight line path.

	1500 byte packet	
Link Rate (Mbps)	Medium Time (μsec)	Interger MTM Weights
11.0	2542	5
5.5	3673	7
2.0	7634	14
1.0	13858	25

Table 5: Medium Times and Weights for 802.11b Transmissions

9. EXPERIMENTAL RESULTS

A number of experiments were conducted to evaluate the performance of the MTM, ETX, and Min Hop route selection metrics. Wireless nodes equipped with 802.11g[36] radios and running on a Linux platform with the Prism54[37] device driver were used in all tests. The cards were configured with RTS/CTS disabled, 5ms Packet Bursting[4] enabled, 1500 byte MTUs, and automatic rate control using all available rates (1 through 54 Mbps).

The experiments were performed in an indoor office environment at the Johns Hopkins University. A single central location was selected as the destination and 15 additional node locations were selected around the building across multiple floors. The 15 node locations had at a minimum, a marginal direct link to the destination. The nodes were arranged such that several multi-hop routes to the destination were available.

Measurements of both link throughput and loss rates were made in order to determine the paths that would be selected by the routing metrics. The Netperf[38] network benchmark tool was used for measuring TCP throughput, and link loss rates were measured using 1500 byte broadcast packets. Broadcast packets are not retransmitted by the MAC layer

9[4]Packet Bursting provides a throughput increase and temporal fairness model similar to the OAR protocol.

Location	Min Hop		ETX		MTM	
	Hops	Mbps	Hops	Mbps	Hops	Mbps
A	1	18.97	1	18.97	1	18.97
B	1	18.50	1	18.50	1	18.50
C	1	17.84	1	17.84	1	17.84
D	1	9.87	1	9.87	1	9.87
E	1	9.72	1	9.72	1	9.72
F	1	8.29	1	8.29	1	8.29
G	1	7.38	1	7.38	2	8.03
H	1	5.27	1	5.27	2	5.82
I	1	3.59	1	3.59	2	4.47
J	1	2.29	1	2.29	2	9.86
K	1	2.07	1	2.07	3	3.30
L	1	0.72	1	0.72	2	9.68
M	1	0.59	2	1.02	3	4.92
N	1	0.43	2	0.93	3	3.82
O	1	0.00	2	0.31	3	5.23

Table 6: End-to-end Path Throughput

so they accurately measure frame loss rates. The data gathered for each link was used to compute which path each routing metric would select from each location. Netperf and static routing were used to measure the end-to-end TCP throughput obtained across each of the selected paths.

Table 6 shows the number of hops taken, and the resulting average throughput obtained from each location using Min Hop, ETX, and MTM selected paths. The locations are sorted in descending order of throughput. Each test location had at least a marginal direct link to the destination, so the min hop metric always selects the direct link as its path, regardless of reliability. As a result, Min Hop only performs well when the direct link is strong (as in locations A through F), and performs particularly poorly when the direct link has low reliability (as in locations M through O). In case O the direct link had the lowest observed reliability delivering only 15% of the broadcast packets. In this case, the TCP throughput tests were unable to complete, so we report the throughput as 0.

The ETX metric avoids using low reliability links by selecting longer paths. This only occurred in a few locations, in our experiments (M through O). If we look at the results from location L we see that the ETX metric selected a single hop path, but was only able to achieve 0.72 Mbps of throughput. This TCP throughput is consistent with the measured 20% loss rate and a 1 Mbps link rate. In order for the ETX metric to select a two hop path, the loss rates in each direction would have to be about 30% or higher (the point at which a lossy one hop path becomes equivalent to a lossless 2 hop path). Our experiments indicate that in multi-rate networks, ETX only provides protection against using links which are on the fringe of the lowest link rate. As a result, ETX is useful for preventing the Grey Zone problem [9], but otherwise operates the same as minimum hop.

The MTM consistently selected the highest throughput path available in the network. The paths selected by the MTM tend to consist of more hops than either of the other metrics. However each selected hop provides high throughput and near zero loss rates. In location L where Min Hop and ETX selected the direct link with moderate loss rates, MTM selected a two hop path and was able to achieve over 13 times the path throughput. In many instances, MTM selected even longer 3 hop paths, and was able to achieve significantly higher throughput then either Min Hop or ETX. These results demonstrate the importance of taking link capacity into account when selecting high throughput routing paths in multi-rate wireless networks.

10. CONCLUSION

In this work we have shown that the minimum hop metric tends to select paths with long slow links. As a result, these paths have low effective throughput and increase total network congestion. In addition, these paths are likely to contain long links that result in low reliability.

We have presented the Medium Time Metric, an improved technique for route selection in multi-rate ad hoc wireless networks. The Medium Time Metric is proportional to the time it takes to transmit a packet on a given link. This metric selects paths that have the highest effective capacity. We have also shown the optimality of this technique under the full interference condition by presenting a formal theoretical model of the attainable throughput of multi-rate ad hoc wireless networks.

Our experimental results with 802.11g radios show that the Medium Time Metric achieves significantly higher throughput then alternative metrics. We observed up to 17 times more end-to-end TCP throughput than with the Min Hop or ETX metrics. Our results both demonstrate the importance of using medium time for selecting high throughput routing paths, and underscore the need for inter-layer communication in order to efficiently and accurately estimate the medium time.

11. REFERENCES

[1] *IEEE 802.11 Wireless Local Area Networks.* http://grouper.ieee.org/groups/802/11/.

[2] B. Awerbuch, D. Holmer, and H. Rubens, "High throughput route selection in multi-rate ad hoc wireless networks," in *Wireless On-Demand Network Systems (WONS)*, January 2004.

[3] C. E. Perkins and E. M. Royer, *Ad hoc Networking*, ch. Ad hoc On-Demand Distance Vector Routing. Addison-Wesley, 2000.

[4] D. B. Johnson, D. A. Maltz, and J. Broch, *DSR: The Dynamic Source Routing Protocol for Multi-Hop Wireless Ad Hoc Networks. in Ad Hoc Networking*, ch. 5, pp. 139–172. Addison-Wesley, 2001.

[5] C. E. Perkins and P. Bhagwat, "Highly dynamic destination-sequenced distance-vector routing (DSDV) for mobile computers," in *ACM SIGCOMM'94 Conference on Communications Architectures, Protocols and Applications*, 1994.

[6] T. Clausen, P. Jacquet, A. Laouiti, P. Muhlethaler, A. Qayyum, and L. Viennot, "Optimized link state routing protocol," in *IEEE INMIC*, (Pakistan), 2001.

[7] D. S. J. De Couto, D. Aguayo, B. A. Chambers, and R. Morris, "Performance of multihop wireless networks: Shortest path is not enough," in *Proceedings of the First Workshop on Hot Topics in Networks (HotNets-I)*, (Princeton, New Jersey), ACM SIGCOMM, October 2002.

[8] R. Dube, C. Rais, K. Wang, and S. Tripathi, "Signal stability based adaptive routing (ssa) for ad hoc mobile networks," February 1997.

[9] H. Lundgren, E. Nordstrom, and C. Tschudin, "Coping with communication grey zones in ieee 802.11b based ad hoc networks," in *WoWMoM 2002*, September 2002.

[10] D. S. J. D. Couto, D. Aguayo, J. Bicket, and R. Morris, "A high-throughput path metric for multi-hop wireless networks," in *9th annual international conference on mobile computing and networking (MobiCom 03)*, September 2003.

[11] D. Aguayo, J. Bicket, S. Biswas, G. Judd, and R. Morris, "Link-level measurements from an 802.11b mesh network," in *ACM SIGCOMM*, August 2004.

[12] R. Draves, J. Padhye, and B. Zill, "Comparison of routing metrics for static multi-hop wireless networks," in *ACM SIGCOMM*, August 2004.

[13] A. Adya, P. Bahl, J. Padhye, A. Wolman, and L. Zhou, "A multi-radio unification protocol for ieee 802.11 wireless networks," in *BroadNets*, 2004.

[14] S. Keshav, "A control-theoretic approach to flow control," in *SIGCOMM*, September 1991.

[15] R. Draves, J. Padhye, and B. Zill, "Routing in multi-radio, multi-hop wireless mesh networks," in *ACM MobiCom*, September 2004.

[16] A. Kamerman and L. Monteban, "WaveLAN-II: A high-performance wireless lan for the unlicensed band," in *Bell Labs Technical Journal*, pp. 118–133, Summer 1997.

[17] A. R. Prasad and H. Moelard, "WaveLAN-II system design note 225: Enhanced data rate control," March 1999.

[18] G. Holland, N. H. Vaidya, and P. Bahl, "A rate-adaptive MAC protocol for multi-hop wireless networks," in *Mobile Computing and Networking*, pp. 236–251, 2001.

[19] B. Sadeghi, V. Kanodia, A. Sabharwal, and E. Knightly, "Opportunistic media access for multirate ad hoc networks," September 2002.

[20] *IEEE Std 802.11b-1999*. http://standards.ieee.org/.

[21] "The network simulator - ns2." http://www.isi.edu/nsnam/ns/.

[22] "Rice networks group." http://www-ece.rice.edu/networks/.

[23] "Orinoco wireless networks." http://www.orinocowireless.com/.

[24] G. Anatasi, E. Borgia, M. Conti, and E. Gregori, "Wi-fi in ad hoc mode: A measurement study," in *the 2nd IEEE International Conference on Pervasive Computing and Communications (PerCom)*, March 2004.

[25] D. Cheung and C. Prettie, "A path loss comparison between the 5 ghz unii band (802.11a) and the 2.4 ghz ism band (802.11b)," in *Intel Labs Report*.

[26] F. A. Tobagi and L. Kleinrock, "Packet switching in radio channels: Part ii - the hidden terminal problem in carrier sense multiple-access and the busy tone solution," in *IEEE Transactions on Communications*, pp. 1417–1433, 1975.

[27] M. Heusse, F. Rousseau, G. Berger-Sabbatel, and A. Duda, "Performance anomaly of 802.11b," in *INFOCOM*, April 2003.

[28] B. Awerbuch, D. Holmer, and H. Rubens, "High throughput route selection in multi-rate ad hoc wireless networks," in *Technical Report Version 1*, December 2002.

[29] K. Jain, J. Padhye, V. Padmanabhan, and L. Qiu, "Impact of interference on multi-hop wireless network performance," in *ACM MobiCom*, September 2003.

[30] T. Leighton, F. Makedon, S. Plotkin, C. Stein, E. Tardos, and S. Tragoudas, "Fast approximation algorithms for multicommodity flow problem," in *Proc. 23rd ACM Symp. on Theory of Computing*, pp. 101–111, May 1991.

[31] J. Aspnes, Y. Azar, A. Fiat, S. Plotkin, and O. Waarts, "On-line machine scheduling with applications to load balancing and virtual circuit routing," in *Proc. 25th ACM Symp. on Theory of Computing*, pp. 623–631, May 1993.

[32] B. Awerbuch, Y. Azar, and S. Plotkin, "Throughput competitive on-line routing.," in *Proc. 34th IEEE Symp. on Found. of Comp. Science*, pp. 32–40, IEEE, Nov. 1993.

[33] S. Lee, S. Banerjee, and B. Bhattacharjee, "The case for a multi-hop wireless local area network," in *Proceedings of The 23rd Conference of the IEEE Communications Society INFOCOM*, March 2004.

[34] B. Awerbuch, D. Holmer, and H. Rubens, "The pulse protocol: Energy efficient infrastructure access," in *Proceedings of The 23rd Conference of the IEEE Communications Society INFOCOM*, March 2004.

[35] J. Li, C. Blake, D. S. J. De Couto, H. I. Lee, and R. Morris, "Capacity of ad hoc wireless networks," in *Proceedings of the 7th ACM International Conference on Mobile Computing and Networking*, (Rome, Italy), pp. 61–69, July 2001.

[36] *IEEE 802.11g-2003*. http://standards.ieee.org/.

[37] *Prism54 Driver*. http://www.prism54.org.

[38] *Netperf*. http://www.netperf.org.