Effects of Multi-rate in Ad Hoc Wireless Networks

Baruch Awerbuch, David Holmer, Herbert Rubens
{baruch, dholmer, herb}@cs.jhu.edu

Technical Report

Abstract— **An ad hoc wireless network is an autonomous self-organizing system of mobile nodes connected by wireless links where nodes not in direct range communicate via intermediate nodes. Modern wireless devices, such as those that implement the 802.11b, 802.11a, 802.11g, and HiperLAN2 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. Our results motivate the need for additional work in this area to fully develop rate aware routing protocols, and stresses the importance of inter-layer communication for optimal performance in wireless networks.**

I. INTRODUCTION

AD HOC wireless networks are self-organizing multi-hop wireless networks where all nodes take part in the process of forwarding packets. A current trend in wireless communication is to enable devices to operate using many different transmission rates. Many current and proposed wireless networking standards have this multi-rate capability. These include the 802.11b, 802.11a, 802.11g, and HiperLAN2 standards. The reason for this multi-rate capability stems directly from 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 no longer 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 the 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.

The rest of the paper is organized as follows. We further define our multi-rate model and simulation environment in Section II. We examine shortcomings of the traditional min hop selection technique in Section III. In order to fully understand the effects of the physical and MAC layers on network throughput, we present a detailed analysis in Section IV. Finally we discuss future research directions and conclude in Section V and Section VI.

Authors are with the Computer Science Department, Johns Hopkins University, Baltimore, Maryland.

II. MULTI-RATE MODEL

The multi-rate model presented in this paper is based on the 802.11b standard [1]. 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 [2] 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 [3] that contain implementations of the RBAR [4] and OAR [5] auto rate protocols. The 802.11 MAC and physical wireless parameters were further modified to match the published specifications of a Lucent ORiNOCO PC Card [6], a commonly used 802.11b wireless adapter (see Table I). Since the carrier sense (CS) threshold specification is not published, we provide an estimate. This estimate was produced by setting the difference between the carrier sense threshold estimate and the 1.0 Mbps receive threshold equal to the difference between the NS2 default carrier sense threshold (-78 dBm) and default receive threshold (-64 dBm).

Figure 1 and Table II show the ranges resulting from these simulation parameters. Real world ranges are smaller due to non-zero system loss, additional noise sources, obstructions, and propagation effects beyond the simple two ray ground model. The results presented here should be valid for any set of ranges with similar proportions regardless of magnitude.

TABLE I
NS2 SIMULATION PARAMETERS

Parameter	Value
Frequency	2.4 GHz
Transmit Power	15 dBm
11.0 Mbps Sensitivity	-82 dBm
5.5 Mbps Sensitivity	-87 dBm
2.0 Mbps Sensitivity	-91 dBm
1.0 Mbps Sensitivity	-94 dBm
Carrier Sense Threshold	-108 dBm
Capture Threshold	10
Propagation Model	Two Ray Ground
System Loss	0 dBm

TABLE II
LINK RATE RANGES AND EFFECTIVE THROUGHPUT

Rate (Mbps)	Maximum Range	Throughput (Mbps)
11.0	399 m	4.55
5.5	531 m	3.17
2.0	669 m	1.54
1.0	796 m	0.85
CS	1783 m	-

III. MINIMUM HOP PATH

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.

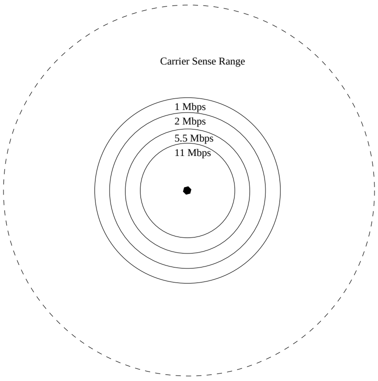

Fig. 1. 802.11b Ranges

A. *Throughput Loss*

In multi-rate wireless networks, the selection of minimum hop paths typically results in paths where the links operate at low rates. 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.

B. *Reliability Loss*

Several recent publications have noted the reliability problems that minimum hop paths are prone to, even in single rate networks, and work has begun on designing protocols that attempt to address this problem [7][8][9]. The origin of this reliability loss, is that the long links in a minimum hop path are near the maximum range a link can be established. As a result these long range links can have very high bit error rates (BER) causing many packets to be undecipherable. A high BER is particularly damaging to most routing protocols, because they use small broadcast packets to establish/maintain routing, but then full RTS/CTS/DATA/ACK exchanges for the real traffic. In this case the small single broadcast has a much higher probability of being delivered than any data traffic, so a routing protocol can be continually fooled into using an unreliable link.

Multi-rate wireless devices are inherently designed to deal with changes in channel quality due to mobility and interference. In 802.11b, as the channel quality degrades, for example when two nodes move in opposite directions, the auto rate protocol will gracefully reduce their link speeds from 11 Mbps down through 1 Mbps before the nodes are finally disconnected. Thus no reliability problems are encountered until reaching the limits of the slowest rate, at which time the exact same high BER problems are encountered in multi-rate networks as in single rate networks. Since a minimum hop path tends to select maximum range links, it can not take advantage of the multi-rate capability of downshifting to a lower rate because links already operate at the lowest rate.

IV. THROUGHPUT PHENOMENA IN MULTI-RATE AD HOC WIRELESS NETWORKS

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.

A. 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.

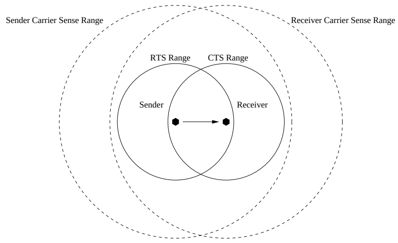

Fig. 2. 802.11 Range

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 of the lack of collision detection hardware in wireless devices. The second is that the 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 back-off 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 II shows the results of a simple NS2 experiment where 1472 byte UDP packets were flooded across a single link. 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 link. In other words, 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

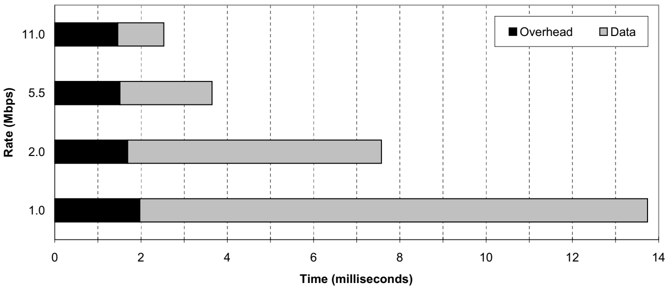

Fig. 3. Medium Times for 802.11b Transmissions

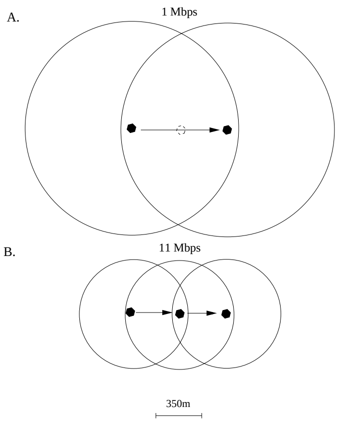

Fig. 4. Path Selection Options

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.

B. 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).

TABLE III

TWO HOP PATH THROUGHPUT

Link Rate (Mbps)		Path Throughput (Mbps)
1^{st} Hop	2^{nd} Hop	
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

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 II and III).

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 11 Mbps. This is because only a single transmission can occur at a time in 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.

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 [10]. 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 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.

C. 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.

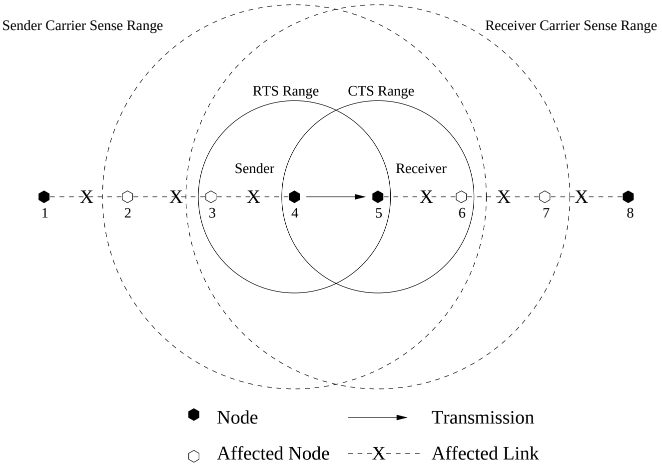

Fig. 5. Effect of Transmission on Other Nodes

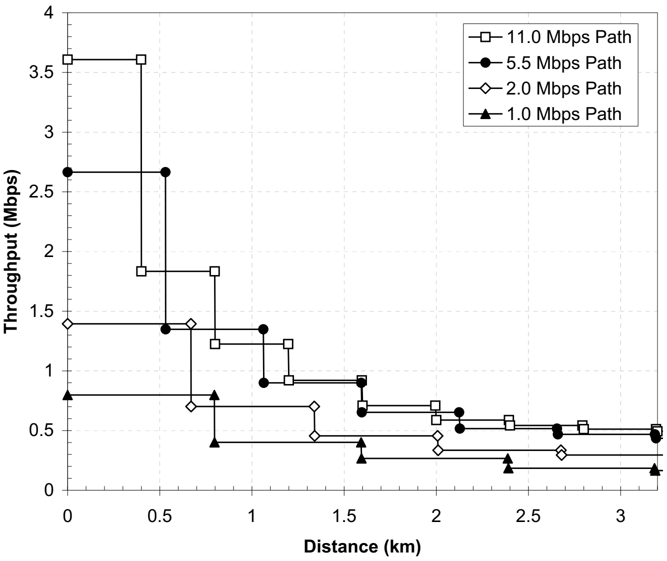

Fig. 6. Throughput Loss Along a Path

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 0.4 km. 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 2.2 km 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.

D. 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

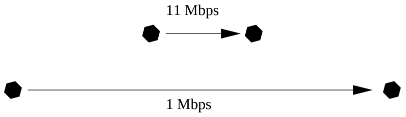

Fig. 7. Temporal Fairness Simulation Node Configuration

TABLE IV
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

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. 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 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 the NS2 network simulator with nodes arranged as indicated in Figure 7. The simulation consisted of two nodes flooding packets to their destination. 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 IV.

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.

V. RESEARCH AGENDA

In order to enable high network throughput in multi-rate networks, a routing protocol must take into account all the above phenomena when making routing decisions. It is the purpose of this paper to bring these multi-rate issues to the attention of the ad hoc wireless networking community in order to foster the development of such routing protocols and routing metrics. Due to the fact that nearly all wireless devices on the market today employ multi-rate capabilities, the development of routing protocols that are capable of extracting the full performance capabilities of these devices is of great importance to the educational, commercial, and government sectors.

In an attempt to take a step forward in this proposed research direction, we have developed the Medium Time Metric (MTM) as described in [11]. The MTM selects paths such that the sum of the medium time consumed by transmitting a packet over each link along the path is minimized. This metric captures many of the phenomena described above, is simple in nature (additive metric), easy to implement in pro-active protocols, shown to be optimal for small networks (with a diameter less than the interference range), and has good performance for larger networks. However much work remains to be done in the development of multi-rate aware routing protocols. Specifically, comparative simulation with modern proactive routing protocols, development of multi-rate aware on-demand protocols, more complete theoretical analysis of larger networks, and the development of new and unique approaches to the multi-rate aware routing problem.

VI. CONCLUSION

In this work we have discussed many of the medium access control and physical layer phenomena that affect routing performance in multi-rate wireless networks. We have shown how these phenomena cause routing decisions to be significantly more complex than in single rate networks, resulting in the traditional single rate hop count metric achieving sub-optimal performance. The results of our analysis show that a throughput increase of up to several hundred percent may be possible by an optimal multi-rate routing protocol. Due to the widespread deployment of multi-rate capable devices, and the importance of throughput in a shared medium environment, we advocate a new thrust by the research community to develop new routing protocols that are multi-rate aware.

REFERENCES

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

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

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

[4] Gavin Holland, Nitin H. Vaidya, and Paramvir Bahl, "A rate-adaptive MAC protocol for multi-hop wireless networks," in *Mobile Computing and Networking*, 2001, pp. 236–251.

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

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

[7] Douglas S. J. De Couto, Daniel Aguayo, Benjamin A. Chambers, and Robert 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, October 2002, ACM SIGCOMM.

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

[9] Henrik Lundgren, Erik Nordstrom, and Christian Tschudin, “Coping with communication grey zones in ieee 802.11b based ad hoc networks,” in *WoWMoM 2002*, September 2002.

[10] 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*, 1975, pp. 1417–1433.

[11] Baruch Awerbuch, David Holmer, and Herbert Rubens, “High Throughput Route Selection in Multi-Rate Ad Hoc Wireless Networks,” Technical Report, www.cnds.jhu.edu/archipelago/.