PARMA: A PHY/MAC Aware Routing Metric for Ad-Hoc Wireless Networks with Multi-Rate Radios*

Suli Zhao and Zhibin Wu
WINLAB, Rutgers Univ.
Piscataway, NJ 08854
{sulizhao, zhibinwu}@winlab.rutgers.edu

Arup Acharya
IBM Research Center
Hawthorne, NY 10532
arup@us.ibm.com

Dipankar Raychaudhuri
WINLAB, Rutgers Univ.
Piscataway, NJ 08854
ray@winlab.rutgers.edu

Abstract

Ad-hoc wireless networks with multi-rate radios (such as 802.11a, b, g) require a new class of MAC/PHY aware metrics that take into account factors such as physical-layer link speed and MAC-layer channel congestion. Conventional "layer 3" ad-hoc routing algorithms typically make routing decisions based on the minimum hop-count (MH). Use of the MH metric leads to selection of paths with few hops but one or more of these hops may turn out to be low-speed radio links due to adaptive rate selection at the physical layer. In this paper, we investigate a new cross-layer routing metric that takes into account both physical layer link speed as well as estimated channel congestion, thus aiming to minimize end-to-end delay that includes both transmission and access times. The proposed "PARMA" routing metric will thus help spread the traffic across the "good links and nodes" in the network, increasing network capacity and reducing packet loss and delay. This paper presents the design and implementation of the proposed PARMA metric for proactive ad-hoc routing protocols such as DSDV. DSDV modifications for incorporating the MAC/PHY aware metric into an ns-2 simulation model are given. Simulation results for typical multi-rate 802.11 ad-hoc network scenarios show that the proposed cross-layer PHY/MAC aware metric achieves significantly higher network throughput and decreases network congestion by selecting paths with high bit-rate links while also avoiding areas of MAC congestion.

1. Introduction

Ad-hoc networks in which radio nodes communicate via multi-hop routing without wired infrastructure have long been considered for military communications. More recently, ad-hoc radio techniques have migrated to commercial scenarios such as sensor networks, home computing and public wireless LAN. Wireless mesh networks are being deployed in cities to provide ubiquitous wireless coverage for the general population or in many instances, as a network for common use by the different first responder agencies such as the police, fire-fighters or emergency medical services.

Ad-hoc networks generally use short-range radios such as 802.11a, b, g [1], Bluetooth [2] or Zigbee [3] as the basic building block. Most of these radios are designed to adapt the physical layer (PHY) bit-rate as a function of channel quality, e.g., the raw transmission speed of an 802.11a radio may vary from 6 Mbps to 54 Mbps depending on received signal strength (RSSI) measurements. Since the PHY bit-rate of each radio link varies with time and node location, it is important to design ad-hoc routing algorithms that take into account dynamic variations in link-speed.

Most conventional ad-hoc routing protocols, including DSDV [7], AODV [13] and DSR [14], use the minimum hop-count (MH) as the metric to make routing decisions. This is primarily a carry-over from wired networks where the transmission rate of a link does not dynamically change and the link rate is independent of the physical transmission range. However, in case of wireless networks, the MH metric tends to choose paths with fewer hops. And each hop in paths chosen by MH will tend to have a longer physical span and also be associated with a lower bit-rate than an alternative path with more number of hops. In order to take advantage of the multi-rate capability and make better use of available network capacity, it is clear that transmission rate needs to be incorporated into the routing metric.

In addition to the transmission rate, we observe that it is also possible to provide an awareness of congestion at each node in order to avoid bottleneck regions with high link utilizations. The wireless link is usually shared with other links in the same neighborhood, while in a wired network, links operate independently of each other and channel access on one link has no effect on any of the adjacent links.

*Research supported by NSF ORBIT Testbed Project, NSF NRT Grant #ANI0335244.

Thus, it makes sense to devise metrics that account for both congestion and rate in a combined manner. For example, a link may provide for a high transmission rate, but could appear congested because neighboring links have a high link utilization. If we account for only the link rate, this link would show up as a “good” link but when combined with a congestion metric, it may turn out to be just the reverse, which is a more accurate reflection of the PHY/MAC layer. Thus, we study both these effects in this paper.

Based on the above considerations, we are investigating a new integrated MAC/routing policy based on a routing metric related to both the PHY bit-rate and MAC congestion information. Taking DSDV as the routing protocol baseline, we study the distance-vector routing behavior under a multi-rate PHY, and with a MAC-aware metric. The problem is studied in detail with an *ns*-2 simulation model [8].

The rest of this paper is organized as follows. In Section 2 we introduce the wireless multi-rate capability and medium access delay. In Section 3 we discuss the routing metrics, and propose a routing metric which integrates the PHY/MAC information such as the link speed and medium busy degree. Then we discuss the potential weakness of distance-vector routing algorithms with the integrated routing metric implemented. We also present enhancements to DSDV by which it can settle down in a reasonably short time and achieve the best routes. The simulation model is presented and the performance is compared with the MH and MTM [4] metrics in Section 4. The last section summarizes the main results and outlines our future work.

2. Multi-rate in MAC Layer

2.1. Overview

The widely used IEEE 802.11x standard uses adaptive selection of physical layer bit-rate as a function of observed channel quality. 802.11b radios can choose different physical rates of 1, 2, 5.5 or 11 Mbps while 802.11a/g radios select between 6, 9, 12, 18, 24, 36, 48 or 54 Mbps. This automatic PHY bit-rate adaptation feature is considered to be useful in most systems because it permits end-users to take advantage of good-quality short-range links when available. When such multi-rate radios are used to build ad-hoc networks, the network topology and link speed changes more dynamically than a single fixed bit-rate is used.

Fig. 1 depicts the way in which an 802.11b radio device experiences different bit-rates when connecting to its neighbors at various distances. If a node wants to use rate 11 Mbps, only nodes in the inner-most circle can decode its frame correctly with sufficient SNR (Signal-to-Noise Ratio). However, if it chooses to use the lower 1 Mbps rate, the transmission range would be much larger. Note that in the above descriptions a simple path loss channel model is assumed and the received signal strength is simply compared to a series of fixed thresholds. Networks may benefit from connections with multiple short-range, high-speed links relative to a single low bit-rate hop that spans a longer distance, as shown in the chain topology of Fig 2.

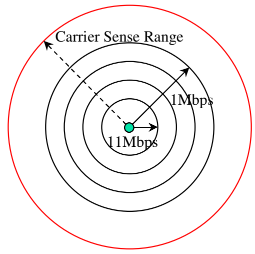

Figure 1. Transmission range of multi-rate.

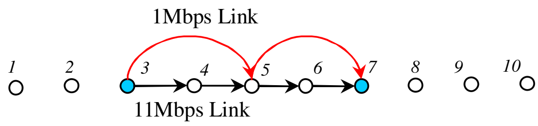

Figure 2. Chain topology and different links.

In Fig 2, 10 stationary nodes are placed in a straight line. Assume each node can reach its immediate neighbor with a fast 11 Mbps link but can only reach the node next to the immediate neighbor with a 1 Mbps link. Therefore, from node 3 to node 7, node 3 can either choose a 4-hop route 3-4-5-6-7 with 11 Mbps rate used for each hop, or a 2-hop route 3-5-7 with 1 Mbps each hop. This diversity of route selection will not occur in single-rate networks. With the simulation in Section 4, it will be demonstrated that short high-speed links are actually better than long-range slow links under certain circumstances.

2.2. Auto Multi-rate Mechanism

The IEEE 802.11 standard does not specify how to choose PHY rate based on varying channel conditions. Previous studies have proposed some schemes to select rate adaptively, such as ARF [9] and RBAR [10].

When constructing an ad-hoc network with muti-rate 802.11x or other similar radios, PHY rate-adaptation can be applied on a packet-by-packet basis depending on the communicating neighbor. This can be implemented with a small table within the MAC layer for recording the selected PHY rate for each neighbor (based on SNR measurements of received packets). Then the device driver can look up the table for the destination of each outgoing frame and obtain a suitable rate for it. For the purpose of our research, we implement this SNR-based autorate scheme in *ns*-2 [8].

2.3. Estimating Channel Access Delay

The channel access delay is an important metric that relates to the offered traffic at the MAC layer. Because the wireless medium is shared, whether a packet could access the channel immediately is determined by not only the states of two ends of the link, but also the states of all neighboring nodes. To measure this effect, a "virtual access delay" estimation is introduced. One intuitive approach is to send periodic probes, but this would introduce extra overhead. We propose a passive estimation method to avoid the probe overhead. Every node records every channel event sensed from physical channel and makes an estimation of the "expected delay if a packet has to be sent". Suppose the nodes in a neighborhood which have packets to access the common channel in a time instant are modeled as in an M/M/1 queue, with delay given by:

$$T_q = T_s \frac{\rho}{1-\rho}$$

where ρ represents the utilization of server (i.e. channel). Each node can estimate ρ, by sensing the occupancy of channel. If it uses an average time cost for one channel event to represent service time T_s of a packet, it could calculate T_q, which is used in the routing metric.

To evaluate the above channel access delay estimation method, we use a grid topology having 154 nodes with spacing of 350 meters between adjacent nodes, as shown in Fig. 3. The carrier sense range is 1783 meters which is slightly greater than 5 hops distance, and other parameters are given in Table 1. We have a 3-hop CBR flow running in the center of the grid. The bit-rate used for each hop is 11 Mbps. The obtained channel occupancy estimates when simulated with saturated load are plotted in Fig. 4. Obviously, the congested area is much larger than the region through which the flows pass in this example. Nodes in the area with the highest busy degree actually lose the ability to support any further flows. The maximum channel busy level is observed to be around 77%.

The access delay estimated by a specific node (shown as the hollow node in the center of Fig. 3), in terms of different offered load, is plotted in Fig. 5. It shows that the estimated delay increases monotonically when the offered load is increased. When the network is congested, the access delay estimation does not increase any further but holds at a steady level. This is because queuing effects cannot be monitored by the PHY layer. Although the queuing delay will increase dramatically if the congestion is not eased, it cannot be estimated from channel access delay alone. We have also compared estimates of access delay with the corresponding measurements and found that the estimate is faithful. The only difference is that estimate of access delay is usually smaller than the actual measurements. There are some additional delays introduced by IEEE 802.11 MAC, such as SIFS, DIFS and backoff intervals. As those delays cannot be monitored, the delay estimation is expected to be an underestimate using the proposed method.

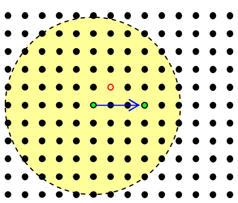

Figure 3. A 14 by 11 grid.

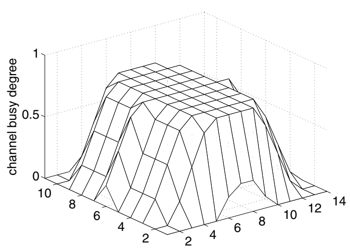

Figure 4. Distribution of channel busy degree.

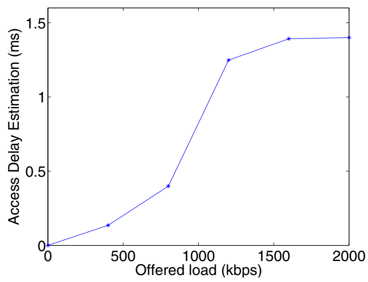

Figure 5. Access delay estimation.

3. Route Selection

3.1. Routing Metric

Previous work has showed that in a wireless ad-hoc network, minimizing the number of hops would choose routes with a small number of links with relatively long physical span and hence lower bit-rate and worse link quality [4, 11]. This motivates investigation of routing metrics which take into account the PHY and MAC layer information.

We propose a routing metric which aims to optimize the packet end-to-end delay. The end-to-end delay of a packet of size L_{pkt} transversing a path p_i is calculated as:

$$Delay_{\,p_i} = \sum_{\forall links \in p_i} (T_{transmit} + T_{access} + T_{queuing}) \quad (1)$$

where $T_{transmit}$ denotes the packet transmission time, T_{access} the medium access time spent by the packet getting access to the link, and $T_{queuing}$ the queuing time required for the packet waiting before trying to access the channel.

The packet transmission time can be calculated as:

$$T_{transmit} = N_{transmit} \times \frac{L_{pkt}}{R_s} \quad (2)$$

where R_s is the link speed, which would be one of the rates the multi-rate devices provide. $N_{transmit}$ is the number of transmissions, including retransmissions, needed for the packet to be received correctly. When the link quality is poor, packet retransmissions will be carried out by the MAC protocol. With adaptive multi-rate PHY, around 90% of packets get transmitted successfully in the first attempt [12]. Hence $N_{transmit}$ can be set to 1 as an approximation.

The medium access time, T_{access}, is used to indicate the medium busy level around the sending node. When the medium is busy, it takes a relatively long time for a packet to get the chance to transmit. Incorporating the medium access time to the routing metric, the routing algorithm can choose a route with light traffic load in addition to high speed links, thus spread the traffic over the network to achieve load balance and increase the effective bandwidth. Note that the access delay estimated by the MAC layer can not be directly use as T_{access}. This will be discussed in Section 3.2.3.

A large access delay reflects a growing interface queue length when the network is congested. When a system below saturation is considered, $T_{queuing}$ can be omitted [1].

With the above simplifications, the routing metric computation can be summarized as Eq.(3).

$$Delay_{\,p_i} = \sum_{\forall links \in p_i} \left(\frac{L_{pkt}}{R_s} + T_{access} \right) \qquad (3)$$

It is clear that this routing metric is both PHY rate-aware and MAC traffic-aware. In the following, we study this kind of routing metric, in conjunction with the class of distance-vector routing algorithms, taking DSDV as a specific example. Our study also reveals potential problems in the cross-layer design of ad-hoc wireless networks.

3.2. Implementation

Our rationale for choosing DSDV is mainly that the periodic DV routing updates can exchange the PHY/MAC information via the metric, so that the dynamic network conditions can be known over the network. Thus nodes can switch their routes whenever there are better routes available. With on-demand routing protocols, routes continue to be used until they are broken. So in order to make on-demand routing protocols work with other metrics than MH, extra control messages have to be added to initiate route discovery procedures before routes are broken. Moreover, the PHY bit-rate can be adjusted according to the received SNRs of the periodic updates. The medium busy degree can also be measured through transmitting DSDV control messages. Probe packets can be used to make measurements as in [5], but extra overhead has to be introduced. With our implementation, extra routing overhead is minimized.

[1]Through the experiments we have verified that the route selection is not affected when the queuing delay is taken into account the metric.

3.2.1 DSDV Operation

DSDV uses a sequence number (SN) which is originated by the destination to indicate the freshness of the routing information and prevent routing loops. In addition, each entry in the routing table is associated with the weighted average settling time, which is the length of time between the arrivals of the first and the best route to a particular destination with the same SN. Since the routing information broadcasts are asynchronous, some fluctuation of routing updates can occur. To solve this problem, and also reduce the number of rebroadcasts, advertisement of routes is delayed until the route has stabilized, i.e., twice the average settling time has passed since the first route is received.

With these two techniques implemented, the best routes are easy to achieve when the MH metric is used. But problems arise when the best route is not the shortest one. This is because, in DSDV, the time for a route from the destination to reach the other end depends on the settling time at each intermediate node in the path. The more the number of hops transversed, the greater the total settling time required. If the best route is not the shortest one, it is likely to arrive late, and even worse, it may arrive later than some routes with a new sequence number. Without the correct settling time, it is difficult to get the best routes.

3.2.2 Enhancement to Achieve Correct Settling Times

As discussed in Section 3.2.1, DSDV uses the delay-advertisement approach. Another approach, called the delay-use, is proposed in [5]. In the delay-use approach, two routing tables are used at each node. A route is not used until it is allowed to be advertised. Before the route can be used and advertised, if there is a route to the same destination with the new sequence number arrived, the old route is moved to the second table and the new route is stored in the current table. We use this modification in our implementation. However, the correct settling time for each entry of the routing tables is critical for both these approaches.

We propose an approach to quickly adjust and achieve the correct settling time. In our approach, the received routes with the last old sequence number are handled instead of being ignored as in the current protocol. This avoids missing the best route when it arrives later than the first route of the next new sequence number. In particular, a route with the last old sequence number is chosen if it has a better metric than the one stored in the second table. Meanwhile, the average settling time of the route to the same destination in the current table is updated accordingly. With this enhancement, the settling time converges quickly and the best routes can be guaranteed before twice the settling time has passed. Our simulations show that 99% of the routes are the best ones with this enhancement, while only one third are the best without this enhancement.

We note that DSDV with two routing tables works well in small and medium networks. When the network becomes large and the variations of the route arrival times increase, the overlap between routes to the same destination but with different sequence numbers will increase. More routing tables are required to store the routing states and prevent missing the best routes. In this sense, DSDV suffers from a scalability limitation when used with PHY/MAC aware metrics.

3.2.3 Enhancement to Smooth Link Layer Changes

Because of the contention access nature of 802.11 MAC, the medium access delay is a random variable. In order to reflect the actual medium busy degree, a smoothing window is used to get the average value over a time duration. Observing the non-linear behavior in Fig. 5, we also propose a non-linear mapping between the estimated access delay and T_{access}. In particular, a time threshold τ is defined according to the network conditions. If the average access delay is less than τ, it is used directly for T_{access}; if it is greater than τ, it is enlarged (e.g., ten times of its value) and then used for T_{access}. If τ is chosen appropriately, the area with high busy degree can be prevented from being chosen to support too much traffic that could cause congestion.

Moreover, due to the different time scales of the network and PHY/MAC layer variations, we must be careful not to degrade the routing performance when incorporating the PHY/MAC aware routing metric. We propose that only significant changes, such as a delay having 20% increase, trigger routing updates; non-significant changes are advertised by periodic updates. Thus, the number of triggered updates are reduced while significant link changes can be advertised over the network. A cautious approach is required for this type of cross-layer design in order to balance factors such as settling time, overhead and routing performance.

4. Simulation Results

The system performance with the proposed PARMA is compared with MH and MTM [4] using *ns*-2 [8]. The performance are evaluated in terms of the system throughput, packet delivery ratio and average end-to-end delay.

4.1. Simulation Parameters

In our simulations with DSDV, the time period between the periodic updates is 15 seconds, the minimum time between the triggered updates is 1 second. An update must be heard from a neighbor in 45 seconds, otherwise the neighbor is regarded as unreachable.

Our simulation study considers constant bit rate (CBR) as the traffic generation model [8]. Packets have a constant size of 512 bytes and are sent at a deterministic rate. The sending rate is varied as an input parameter to gradually increase the offered load to the network.

Multi-rate 802.11b is used. The transmission power is fixed at 15 dbm. RTS/CTS is disabled [2]. ACKs are transmitted using the basic rate of 1 Mbps. The receiver thresholds and the corresponding effective distances are shown in Table 1, which are obtained by using the two-ray ground reflection propagation model [8].

Table 1. Simulation parameters.

	Receiver threshold	Distance
Carrier sense	-108 dbm	1783 m
1 Mbps rate	-94 dbm	796 m
2 Mbps rate	-91 dbm	669 m
5.5 Mbps rate	-87 dbm	532 m
11 Mbps rate	-82 dbm	399 m

4.2. Scenario I

First we try a simple topology with 10 stationary nodes placed in a straight line with a distance of 350 meters between the neighboring nodes, as shown in Fig. 2. In this linear scenario, a short link with distance of 350 meters can provide 11 Mbps data rate, while a longer link with distance of 700 meters only gives 1 Mbps. There is one flow, originated from one node to the destination which is 1400 meters away. As observed, the MH metric tends to choose the shortest path with two 1 Mbps links, and the PARMA metric chooses the route with four 11 Mbps links. The simulation results are shown in Fig 6 and 7.

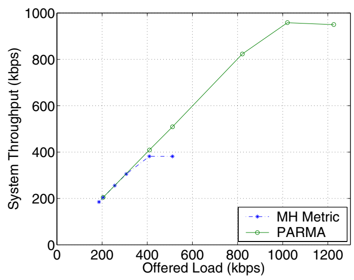

Figure 6. Throughput vs offered load.

The results show that in the chain topology, system performance improves with the PARMA metric. In this scenario, there is a significant factor of 2.5x improvement in system throughput. Also, the packet delivery ratio and the average end-to-end delay are improved. In this scenario, PARMA has the same behavior as MTM.

[2]The experiment results with RTS/CTS enabled are similar to those with RTS/CTS disabled, though with slightly lower improvements. This is because the basic rate is used for sending the RTS and CTS frames, and this would increase the MAC overhead, especially for high rate links.

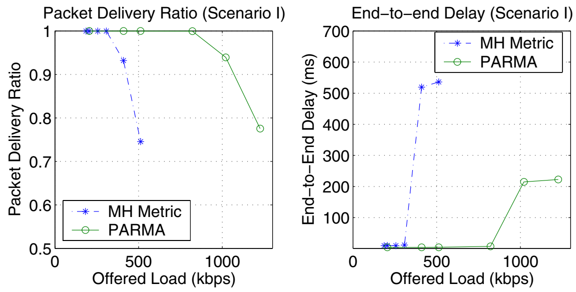

Figure 7. Packet delivery ratio and delay.

4.3. Scenario II

A 6 by 7 regular grid topology is used as the second evaluation scenario. The distances between neighboring nodes in both the horizontal and vertical directions are 350 meters. There are three possible rates in the topology: 1, 5.5 and 11 Mbps. All nodes are stationary. There are two traffic flows in the network: flow 1 is from node 2 to node 3, and flow 2 is from node 30 to node 35, as shown in Fig. 8.

We choose flow 1 to start transmitting packets earlier than flow 2. The traffic generation rate of flow 1 is fixed to be around 3 Mbps in order to keep the medium around flow 1 busy enough such that the congestion can occur when flow 2 is added; and the traffic generation rate of flow 2 is varied as the input parameter to gradually increase the offered load. As shown in Fig. 8, node 2 and node 32 are in the carrier sense range of each other, so there is interference between these two flows if flow 2 takes a route including the link from node 32 to 33. Then flow 1 and flow 2 will compete for the medium when both have started. If flow 1 is close to saturation and there is a congested area around it, an ideal routing protocol will guide flow 2 to go around this congested area and achieve load balancing that prevents the whole system from becoming congested.

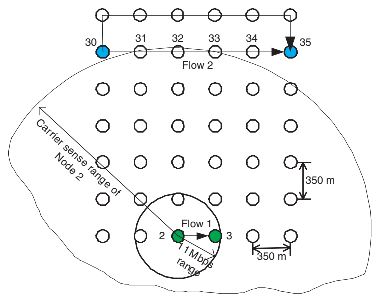

Figure 8. Grid topology and flows.

We study how the system handles interfered flows and avoids the congested area when using different routing metrics. Fig. 9 shows how the total throughput changes with time. In each run, flow 1 and flow 2 start at 50 sec and 150 sec, respectively. Fig. 9(a) and (b) imply that broken links exist due to interference, which results in packet drops. Fig. 9(c) shows that there are metric changes when flow 2 starts transmissions. After some period of settling, a new route which avoids interference from flow 1 has been used (an alternative route is shown in Fig. 8) and thereafter the system throughput starts to go up. Fig. 9(c) thus reflects the efficacy of our algorithm under dynamic traffic conditions.

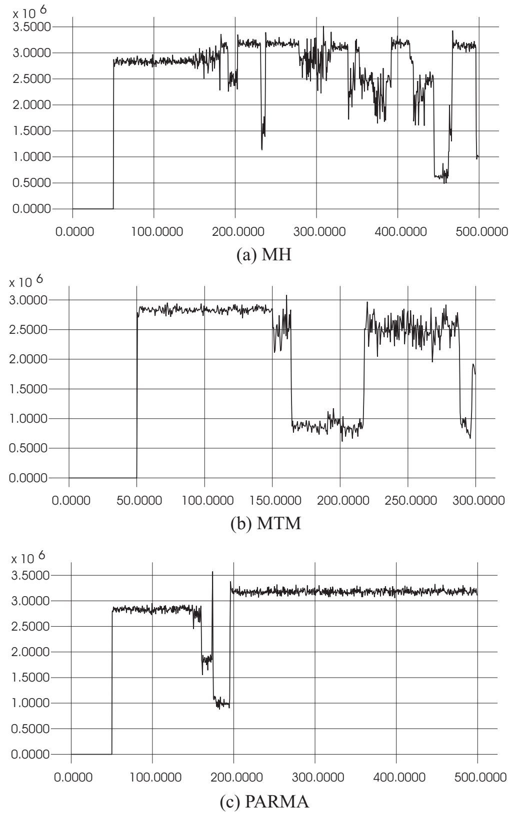

Figure 9. Total throughput vs time.

Fig. 10 gives the throughput achieved using different routing metrics. The x-axis indicates the offered load of flow 2. We observe that using MTM, system throughput drops when the offered load is increased to 550 kbps. System throughput also drops for MH. However, the PARMA throughput curve keeps increasing slowly, and is expected to increase until the system becomes excessively congested. For flow 2, MH switches among the possible 3-hop routes, and MTM always chooses horizontal paths from node 30 to 35. If there is no interference, PARMA will work like MTM. But due to the potential interference, PARMA will guide flow 2 to avoid the congested area, using routes with more hops or low rate links (thus the PARMA throughput of flow 2 is lower than that of MTM).

The end-to-end delay is not plotted due to space limitations. The results show that PARMA can achieve the end-to-end delay comparable to MTM. But when the system is

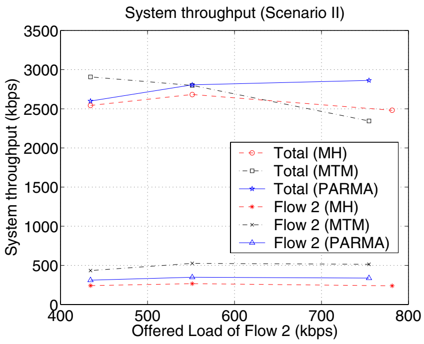

Figure 10. System throughput vs offered load.

close to saturation, PARMA has slightly larger delay. This is because PARMA helps improve the packet delivery ratio thus some packets finally arrive at their destinations but with longer latency than the average value.

5. Related Work

Routing algorithms incorporating multi-rate PHY has been investigated in [4], which proposed Medium Time Metric (MTM) to find paths such that the total transmission time is minimized. MTM is a static solution which only handles the multi-rate transmissions and leaves out the wireless medium access contention factor. A "full interference" assumption is used in its theoretical model.

De Couto et al. [11] observed that using the shortest path would result in poor throughput. And they proposed the expected transmission count metric (ETX) to incorporate the effects of link loss ratios [5]. ETX introduces extra routing overhead of the dedicated link probe packets to measure the delivery ratios. Also ETX is independent of network load and does not attempt to route around congested links.

MAC layer utilization information has been used in [6]. The MAC layer utilization is averaged over a 10-second period and used to adjust the behavior rather than metric of routing protocols. Here we propose to estimate the medium access delay via short-term measurements, and incorporate this information into the routing metric to avoid the busy area without higher layer congestion avoidance schemes.

6. Conclusions and Future Work

We have studied a specific PHY/MAC aware routing metric working with distance vector routing protocols such as DSDV. In order to make DSDV work well with the PHY/MAC aware routing metric, we propose specific enhancements to the routing protocol. In addition, smoothing techniques for the link portion of the proposed metric is introduced to adjust the different change variations between the MAC layer and the network layer and also to improve route convergence. Our simulation results show that routing metrics which only consider the number of hops may not achieve high throughput in multi-rate networks. Using a metric only based on the medium transmission time would have the effect of guiding traffic to high speed links but this could cause MAC layer congestion in some area. With both the PHY rate and the MAC occupancy level taken into account in the routing metric, packets can choose the high rate links while also avoiding congested areas in the network.

Further work is planned on investigating trade-offs between settling time, routing overhead and network performance over the range of possible methods and parameters. Moreover, we are going to investigate PARMA with different routing schemes such as DSR, mobility scenarios, and other traffic patterns. We believe that our results show promising improvements with cross-layer routing metrics (such as PARMA proposed here) but that careful design is needed to avoid unintended effects.

References

[1] IEEE Std 802.11, *Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications*, 1999.

[2] Bluetooth Core Specification v2.0, Bluetooth Special Interest Group, http://www.bluetooth.org

[3] IEEE Draft Std 802.15.4, *Wireless Medium Access Control (MAC) and Physical Layer (PHY) Specifications for Low Rate Wireless Personal Area Networks (LR-WPANs)*, 2003.

[4] B. Awerbuch, D. Holmer, and H. Rubens, "High Throughput Route Selection in Multi-Rate Ad Hoc Wireless Networks", *Proc. WONS 2004*, 2004, pp. 251-268.

[5] D.S.J. De Couto, D. Aguayo, B.A. Chambers, and R. Morris, "A High-Throughput Path Metric for Multi-hop Wireless Routing", *ACM MobiCom 2003*, San Diego, Sept. 2003.

[6] Y.-C. Hu and D.B. Johnson, "Exploiting Congestion Information in Network and Higher Layer Protocols in Multihop Wireless Ad Hoc Networks", *ICDCS 2004*, 2004, pp. 301-310.

[7] C.E. Perkins and P. Bhagwat, "Highly Dynamic Destination-Sequenced Distance-Vector Routing (DSDV) for Mobile Computers", *Proc. ACM SIGCOMM'94 Conference on Communications Architectures, Protocols and Applications*, 1994, pp. 234-244.

[8] K. Fall and K. Varadhan, Eds., "The *ns* Manual (2002)", The VINT Project, http://www.isi.edu/nsnam/ns/doc/ns_doc.pdf

[9] A. Kamerman and L. Monteban, "WaveLAN-II: A High-Performance Wireless LAN for the Unlicensed Band", *Bell Labs Technical Journal*, 1997, pp. 118-133.

[10] G. Holland, N. Vaidya, and P. Bahl, "A Rate-Adaptive MAC Protocol for Multi-Hop Wireless Networks", *Proc. ACM MobiCom*, 2001.

[11] D.S.J. De Couto, D. Aguayo, B.A. Chambers, and R. Morris, "Performance of Multihop Wireless Networks: Shortest Path is not Enough", *Proc. HotNets-I*, Princeton, New Jersey, October 2002.

[12] P. Gopalakrishnan, "Methods for Predicting the Throughput Characteristics of Rate-Adaptive Wireless LANs", (Master thesis, Rutgers University, 2004), pp. 28-32.

[13] C.E. Perkins and E.M. Royer, "Ad Hoc On-Demand Distance Vector Routing", *Proc. IEEE Workshop on Mobile Computing Systems and Applications*, 1999, pp. 90-100.

[14] D.B. Johnson and D.A. Maltz, "Dynamic Source Routing in Ad Hoc Wireless Networks", *Mobile Computing*, T. Imielinski and H. Korth, Eds., Kluwer Publishers, ch. 5, pp. 153-181, 1996.