Experimental Anatomy of Packet Losses in Wireless Mesh Networks

Hua Yu Daniel Wu Prasant Mohapatra

Department of Computer Science
University of California, Davis
Email: {huayu, danwu, pmohapatra}@ucdavis.edu

***Abstract*—Despite the increasing number of wireless mesh network deployments and research, there is a lack of understanding on how robust these networks are in practice. In this paper, we perform a systematic experimental study to investigate the impacts of different factors on the unreliability of a mesh network and the sources causing such unreliability. We use packet loss rate as a metric for defining reliability. The factors that we studied include traffic load, number of hops and flows, transmission rates, maximum retransmission limits and the RTS/CTS mechanism. Our results are based on measurements performed on our real-world mesh network testbed. In addition to performing the experiments on multiple hops, we include the results of one hop experiments for comparison. In identifying the sources of packet loss, we developed a tool, FlowPaC, to collect flow-based statistics at different points in the system to understand the effects of the MAC layer parameters and the traffic attributes. We also explore the potential remedies for the system configuration and thereby improve the reliability of wireless mesh networks.**

I. Introduction

Wireless mesh networks (WMNs) have become increasingly popular for extending last-mile Internet access due to low-cost, off-the-shelf equipment. Many important civilian applications, including, but not limited to, broadband access in residential communities, disaster relief and intelligent transportation systems are popular application environments. Beyond these applications, WMNs are very useful in military deployments. All of these scenarios need a reliable, robust and resilient network. Although a plethora of work is dedicated to performance and capacity of these networks, efforts on enhancing reliability are lacking.

Unlike wired networks, wireless mesh networks also have to endure with open-air shared communication medium, dynamic and time-varying operating environment, and limited resource availability. Consequently, it is desirable and necessary to develop efficient technologies capable of enduring and recovering from any adverse conditions. However, there is a lack of understanding of how mesh networks perform in practice, especially in terms of reliability.

In this work, we perform a systematic study to investigate the effects of different factors on the unreliability of mesh networks and identify the sources that cause such unreliability. To our knowledge, our work is one of the first to systematically and experimentally address the reliability problem in wireless mesh networks.

We define **reliability** from the network user's point of view. It is the acceptable user experience under network component failures in dynamic conditions, i.e., no disruption of service and low packet loss ratio. In this work, we use the packet loss rate as a metric for reliability and investigate the causes of packet loss in the network. Previous work mainly consider the packet loss due to random noise or collisions from interference for a single hop. However, in real networks, the packet loss sources are more than that. In the following section, we first describe the causes of packet loss. Moreover, it is still not clear how the different MAC layer configurations and traffic attributes affect the packet loss, especially for multi-hop flows. Thus we develop a tool, FlowPaC, to collect flow-based statistics from the various points in the system to diagnose the source of packet loss. We find that traffic attributes are important factors in determining the packet losses and there exists complicated dependencies in setting MAC layer parameters.

A. Wireless Packet Losses

There are various causes for packet losses in wireless multi-hop networks. We classify the various reasons into three groups. The first group is **channel induced factors**. This includes the random bit error from signal attenuation, multi-path fading, shadowing and noise. Whether the signal is valid at the receiver largely depends on the received signal strength (RSS) from the transmitter, and that of the interference at the receiver. Given the transmission power, the RSS is mostly decided by path loss over the transmitter-receiver distance, which models the signal attenuation along the distance. Other factors include multi-path fading,shadowing and noise, which have less influence in the open space environment, but can not be ignored in urban environments with many buildings.

The second group is **interference induced factors**, including interfering nodes in or out of the mesh network that operate on the same channel as the desired transmission. In the presence of interfering transmissions, a collision can occur at the receiver when interfering signals corrupt the original transmission and lead to an collision-induced packet loss. Note that if the interfering transmission can be detected within the carrier sensing range, random wait is usually used to avoid the collision of the current packet at the sacrifice of the longer contention time, and hence the reduced effective transmission

978-1-4244-2908-0/09/$25.00 ©2009 IEEE

rate and possible drops of later packets. This is usually the case for a multi-hop flow or multiple flows.

Failures may happen because of packet dropping due to buffer overflow or link degradations. Furthermore, even when the signal arriving at the receiver is valid, packet loss can still happen when the receiving node fails, moves, or simply drops packets because of limited buffer size or slow CPU processing speed. We classify this third group as **node induced factors**. For example, on Linux system, each processor has a `softnet_data` structure, which holds a list with the incoming packets received by the network interface card (NIC). In this structure, the length of the backlog queue is stored. This queue will build up in size when an interface receives packets faster than the kernel can process them. If this queue is too small (default is 300), packets will begin to lose at the receiver, rather than on the air. Moreover, if the size of the socket receive-buffer is too small, and if the application can't keep up with the network speed, there will be more contention for CPU resources and buffer spaces and packets will drop.

B. Contributions and Inferences

- We performed a systematic experimental study to investigate the effects of different factors on the unreliability of the mesh network. To identify the sources causing such unreliability, we developed a tool, FlowPaC, to collect flow-based statistics at various places in a system. Our results are based on measurements performed in our real-world mesh network testbed.
- The high level factors that we studied include traffic load, the number of hops, and flows. We found the following results: Even with the same load, higher packet generation rate results in much heavier packet losses than higher packet sizes; Packet generation rates greater than 1000 pkts/sec show a large amount of packet drops in the CPU's receive queue which makes the difference among the packet sizes less important. In general, the packet loss increases with the number of hops, but if the application cannot keep the network speed and the buffer size is limited, it presents a higher packet loss even for one hop networks. The loss from interference is the main reason for the multi-hop flow packet losses under the light to moderate load. The loss rate is the same for the same load irrespective of the number of flows under the same packet size.
- The MAC parameters that we studied include transmission rate, maximum retransmission limits and RTS/CTS mechanism. In general, automatic rate adaptation with multi-retry mechanism or a higher transmission rate with a small number of retransmission limits is better.
- In identifying the sources of packet loss, we also analyze the loss variability across time. There exists bursts of high packet losses from the medium in light load scenarios, which results in an overall high loss ratio. For high load scenarios, the packet loss starts a few seconds from the beginning of the transmission and almost persists during the rest of transmission as a large amount of packets are consistently dropped after they were received from the medium and before they were put in to the CPU's receive queue for processing.
- Having identified the primary source of losses in mesh networks, we show the potential remedies to mitigate theses losses: increasing the size of the transmission queue, backlog queue and socket receive buffer.

II. Related Works

Aguayo et al. [1] identified interference as a main source of packet loss for a outdoor 802.11 mesh deployment. [2] characterized the packet losses for a WiFi-Based Long Distance Network. A large number of measurement based studies have also been carried out to study the source of packet loss in indoor large scale 802.11 deployments [3]. Most the experiment work only consider the factors from the network. We focused on evaluating the impact of different factors from both the network and the system on the network reliability.

Many approaches have been proposed to improve the reliability of multi-hop wireless networks. At the physical layer, channel coding or multiple antennas [4], [5] can be utilized. At the link layer, Automatic Repeat reQuest (ARQ), rate adaptation/power control [6], [7] and multiple radios [8] can improve reliability. At the network layer, multi-path routing [9] and opportunistic forwarding [10] can provide better reliable paths. At the transport layer, employing more reliable transport protocols [11] can reduce perceived application layer packet losses. However, most of the solutions are based on simulations or attained theoretically.

III. Experimental Methodology

A. Topology and Configuration

We perform our packet loss characterization measurements on a mesh network testbed comprising of four mesh nodes (`node1` to `node4`) with one client (`client1`) shown in Figure 1. Figures will use the short hand notation of `n1`,...,`n4` to denote nodes and `c1` to denote the client. The mesh nodes are situated within a house and placed to increase the wireless coverage.

Each mesh node is a 266 MHz x86 Soekris net4826 embedded device running a custom built Linux distribution using a 2.6.23 Linux kernel [12]. The mesh nodes have 128MB SDRAM main memory and 64MB compact flash for the OS and other storage. The client is a HP nc6000 laptop running with a Linux kernel of 2.6.25. We use Atheros 802.11 a/b/g radios in all of our devices.

Each of the mesh nodes in our testbed is equipped with two radios. One radio is configured to be in access point mode and serves as an access gateway to the network for wireless clients. The other radio is configured to be in ad-hoc mode and participates in routing using the OLSR routing protocol. This radio connects with the wireless mesh back-haul and is configured on the same channel as those of the backbone radios on other nodes. The AP mode radio is configured to channel 48 while the ad-hoc mode radio is set for channel 36

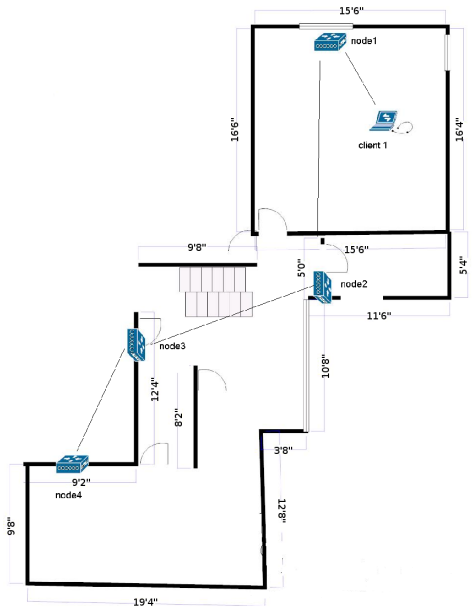

Fig. 1. Experiment Topology

to minimize cross radio interference on the 802.11a channels. In our experiments, `client1` connects to `node1` through the access radio on channel 48. The traffic will be forwarded from `node1` to `node2` to `node3` and finally come to `node4` on Channel 36.

B. Software and FlowPaC Tool

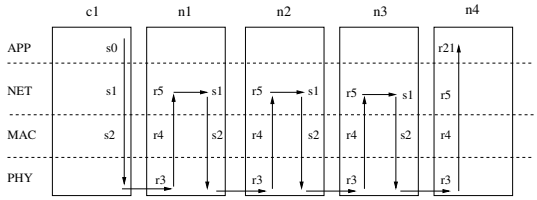

Fig. 2. Data Collection Points

All devices use the madwifi-ng device driver [13] for the radios. We use thrulay [14] to measure packet loss rate.

In order to examine where the packet losses occur and how many packets are lost on its way from the source to the destination node, we developed a tool FlowPaC, which is based on a modification of tcpdump-3.9.8 [15], libpcap-0.9.8, madwifi-ng device driver and the 2.6.23 Linux kernel.

With the FlowPaC tool, we can use the same user interface as Tcpdump to expose the following statistics for a selected flow. Figure 2 denotes where in the system the points are collected from.

1) *s1:FromNet* records the number of packets that are successfully routed from an application or forwarded traffic, and are going to placed into the transmission queue.
2) *s2:XmitedAir* records the number of packets that are actually sent to the medium.
3) *r3:RcvdAir* records the number of packets that are actually received by the device driver from the medium.
4) *r4:Rcvd80211* records the number of packets that are accepted by the 802.11 protocol.
5) *r5:ToNet* records the number of packets that have put into the CPU's receive queue and are going to be forwarded or received by the local host.
6) *s6:XmitDriverFail* records the number of packets that have been placed into the transmission queue but failed due to the device driver. This statistics does not include the packets that are transmitted but failed to reach the next hop.

According to the above statistics, together with *s0:UdpSend* and *r21:UdpRecv* provided by thrulay, we can know where the packet losses occur. More specifically,

- The difference between *s0:UdpSend* and *s1:FromNet* indicates the number of packets that failed due to upper-link-layer errors, such as routing or error packets, for application-generated traffic.
- The difference between *s1:FromNet* and *s2:XmitedAir* indicates the number of packets that failed due to the packet drop in the transmission queue and the failure in the device driver *s6:XmitDriverFail*. One tricky thing here to note is that the madwifi driver re-queues the packets into the transmission queue if the device is temporarily down or the buffer of the driver is taken up. However, the re-queued packets will finally be dropped by the transmission queue, fail in the driver or get a chance to transmit to the medium. In all of our experiments, *s6:XmitDriverFail* is equal to zero, so we didn't include this statistic in the results. The channel-induce, node-induced and interference-induced factors all can cause such packet loss. Here the interference-induced factors leads to reduced effective transmission rate, longer packet in queue time, and packet drop for arriving packets.
- The difference between *s2:XmitedAir* and *r3:RcvdAir* indicates the number of packets that were lost because of collision or random loss on the link. Both the channel-induced and interference-induced factors can cause such packet losses. Here interference-induced factors leads to collision-based packet loss.
- The difference between *r3:RcvdAir* and *s4:Rcvd80211* indicates the number of packets that were lost due to 802.11 protocol semantics. Packet losses for this reason seldom happen according to our experiment results.
- The difference between *r4:Rcvd80211* and *r5:ToNet* indicates the number of packets that were lost due to the packet drop at the CPU's receive queue. The node-induced factors can cause such packet loss.
- The difference between *r5:ToNet* and *s1:FromNet* indicates the number of packets that failed due to upper-link-layer errors, such as routing, for forwarded traffic.
- The difference between *r5:ToNet* and *r21:UdpRecv* indicates the number of packets that failed due to the socket receive queue at the destination host. The node-induced

factors can cause such packet loss.

IV. High Level Impacts on Reliability

In this section, we investigate the non-MAC layer-factors as potential sources of packet loss. All of our results are based on a 20-second constant bit rate (CBR) UDP traffic stream, and the loss rate is an average of five experiments. Unless otherwise indicated, this set of experiment uses a transmission power of 5dBm, 54Mbps transmission rate, zero retries in MAC layer retransmissions and with RTS/CTS disabled for the MAC layer settings. The distance between the source and destination is 4 hops. A traffic load of 8Mbps with CBR packets of 1024 bytes each and 1000 packets per second (pkts/sec) is used. Based on the measurements performed on our mesh testbed, we show three key results:

- Even with the same traffic load, higher packet generation rate results in much heavier packet losses than larger packet sizes. The main point of loss is at the CPU's receive queue.
- Packet loss increases with the number of hops. With more hops, more interference increases the medium loss. The forwarding processing for multi-hop flows further increases the packet drop at CPU receive queue. If the application cannot keep up with the network speed and the buffer size is limited, there presents a high packet loss even for a single hop.
- The loss rate is the same for the same traffic load irrespective of the number of flows using the same packet size.

A. Impact of Packet Generation Rate and Packet Size

Figure 3(a) shows the loss rate by varying the packet generation rate from 500 pkts/sec to 2000 pkts/sec and packet sizes from 50B to 1408B. For each of the packet generation rates, the packet loss increases as larger packets are used.

For the traffic load of 4.1Mbps, at 500 pkts/sec and packet size of 1024B the loss rate is about 20%; at 1000 pkts/sec and 512B the loss rate is 30%; at 1500 pkts/sec and 341B the loss rate is approximately 55%; while at 2000 pkts/sec and 256B is about 70%. Even with a light load of 800kbps, a high packet loss rate of 67% is seen at 2000 pkts/sec with a packet size of 50B while an almost zero loss rate is seen at 500 pkts/sec at the same packet size. We can see that higher generation rate leads to much heavier packet losses compared to larger packet size for the same load.

With the same packet generation rate, the difference of loss rates among all the different packet sizes become smaller as seen from the smaller slope for the higher generation rate in Figure 3(a). The loss rate difference for 2000 pkts/sec is 12% from using 50B to 1408B, while it is 25% for 1500 pkts/sec and 48% for 1000 pkts/sec. We can conclude that packet sizes have less impact on packet losses for high packet generation rates than low packet generation rates.

To find out the source of the high packet loss under high traffic load, we focused on seven different traffic loads by varying the packet generation rates and packet sizes in incremental steps. For each load, we run FlowPaC and record the number of packets in Figure 3(c). The x-axis value corresponds to the statistics order in the previous section. We found that there is 80% and 60% drop of packets from point *r4* to *r5* at `node1` for packet generation rate of 2000 pkts/sec and 1500 pkts/sec, 13% for 1000 pkts/sec and 0% for 500 pkts/sec, which indicates the packet drop at the CPU's receive queue. After the first hop, the high load is reduced considerably, even lower than that of the lower traffic at this point, and there is only 2-8% of medium loss before reaching the destination according to the difference between point *s2* and *r3* at each hop. However, this medium loss constitutes the main loss for small load. Figure 3(c) also shows that packet size imposes little effect on this loss for packet generation rates greater than 1500 pkts/sec as the queue drops the whole packet. For comparison, we show in Figure 3(b) the packet loss rate under the same set of loads within one hop from `node1` to `node2`. The trend is similar to the four hop scenario, only with smaller packet loss rates as there is fewer medium loss from the interference.

To further confirm this finding, we decode the packet loss file generated by d-ITG [16] traffic generator and show the number of packet loss with time in different time intervals, i.e., varying from 200ms to 5s, in Figure 4. Each point denotes the number of packet losses during the corresponding time interval starting from the same time. With packet generation rate of 1000 pkts/sec, there is a burst of packet loss at time 20s and a small packet loss before time 18 seconds in the experiment due to the medium loss. For higher packet generation rates, the

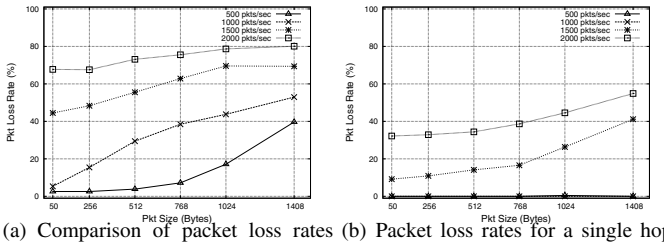

(a) Comparison of packet loss rates for various packet generation rates and packet sizes over four hops.

(b) Packet loss rates for a single hop with various packet generation rates and packet sizes.

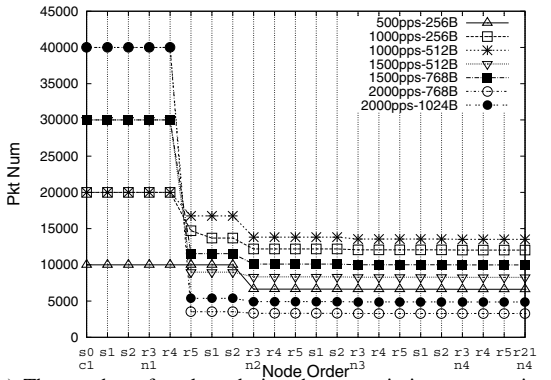

(c) The number of packets during the transmission or reception at each node.

Fig. 3. Experimental results for correlation between packet generation rates and packet sizes.

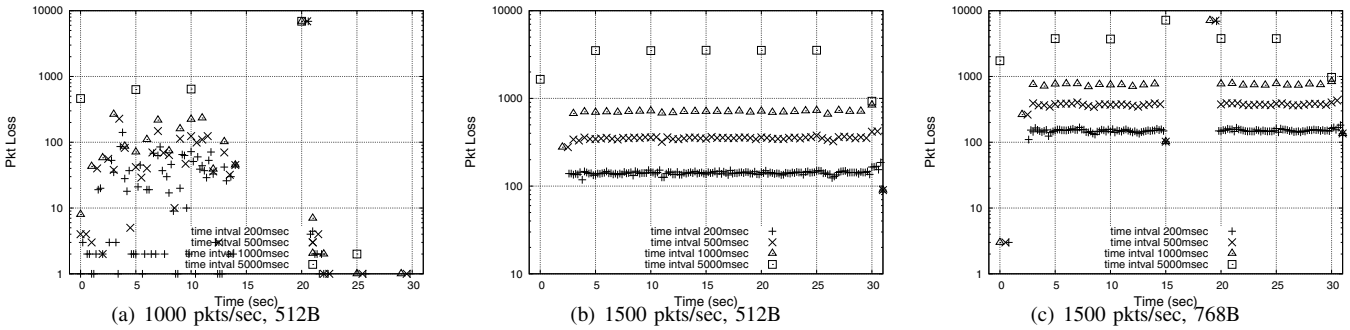

(a) 1000 pkts/sec, 512B (b) 1500 pkts/sec, 512B (c) 1500 pkts/sec, 768B

Fig. 4. Number of packet losses over time at different packet generation rates and packet sizes.

packet loss persists throughout the whole experiment due to the queue drop. This also infers that the effect of the wireless network environment is much smaller compared to the effect of the high traffic load.

B. Impact of Hop Counts and Traffic Load

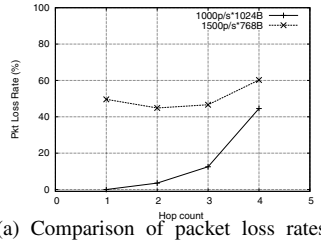

(a) Comparison of packet loss rates for different number of hops under similar loads.

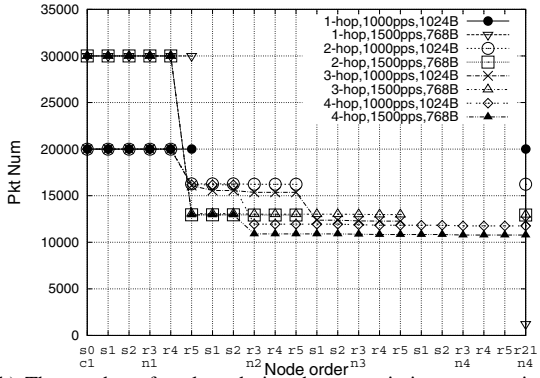

(b) The number of packets during the transmission or reception at each node.

Fig. 5. Experimental results for correlation between packet generation rates and packet sizes.

Figure 5(a) shows the loss rate by varying the number of hops from one to four under different loads. With a moderate load of 8Mbps at 1000 pkts/sec and packet size of 1024B, the loss rate increases with an increase of the number of hops. There is a high increase in loss rate going from 3 hops to 4 hops, a moderate increases from 2 hops to 3 hops, and a small increase from 1 hop to 2 hops. However, with a similar amount of load of 9Mbps at 1500 pkts/sec and 768B packet size, the loss rate decreases from 1 to 2 hops, very few increase from 2 hops to 3 hops, and a moderate increase from 3 to 4 hops. The increase from 3 to 4 hops is smaller for 1500 pkts/sec compared to the case with a smaller rate of 1000 pkts/sec, which is consistent with the result we see from Section IV-A.

To explore this phenomenon, we run our tool, FlowPaC, and record the number of packets at each place in Figure 5(b). We found that all the packets reach the first hop at `node1` from the medium.With one hop, the packets safely reach the destination node until place *n1:r5* and are to be processed by the network layer,but most packets of the flow of 1500 pkts/sec are dropped due to the limited socket receive size, slow CPU processing speed and frequent interrupt from hardware to receive packets at high arrival rate. This result is also confirmed using *netstat* command. For multiple hops, the main packet losses (approx. 53%) is from *n1:r4* to *n1:r5* for 1500 pkts/sec, due to the high packet arrival rate, low processing speed and small CPU receive queue size. The forwarding process for multi-hop flows increases this packet drop. With 4 hops, the increased interference causes the packet losses at the medium, from *n1:s2* to *n2:r3*, which is the main portion for the 1000 pkts/sec flow, and further increases the packet loss rate for 1500 pkts/sec flow. After the packets passed the second hop, the traffic load is reduced considerably. Less packets are lost at point from *n2:r3* onwards.

Under a similar load, although the packets were received by the network device at the first hop, the packet loss at the CPU's receive queue is of considerable large amount for high packet generation rate flows. This is even higher than the loss due to the bad network condition. For multi-hop flows, the forwarding process increases the packet drop at CPU receive queue and the interference increases the medium loss.

C. Impact of Flow Counts and Packet Generation Rate

Figure 6(a) shows the loss rate by varying the number of flows from one to six under the automatic rate control and 54Mbps rate. Each basic flow is a 2Mbps CBR traffic of 512B each at a packet generation rate of 500 pkts/sec. All flows start at `client1` and end at `node4`. The packet loss rate increases with an increase in the number of flows. This trend is the same for both auto rate and 54Mbps transmission rate.

To study the relationship between the number of basic flows and the packet generation rate, we compare the loss rate of x

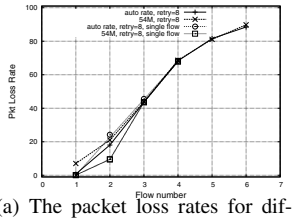

(a) The packet loss rates for different number of flows and packet generation rate.

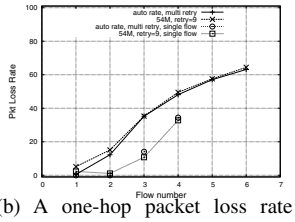

(b) A one-hop packet loss rate for different number of flows and packet generation rate.

Fig. 6. Experimental results for correlation between the number of flows and packet generation rate.

number of basic flows with the loss rate of a single flow under the generation rate of x times that of basic flow 500 pkts/sec and the same packet size 512 bytes. We found that for both auto rate and 54Mbps transmission rate, the loss rate of a single flow under the generation rate of x times of 500 pkts/sec is equal to that of x number of basic flows 500 pkts/sec each. Also, the loss rate of each basic flow is similar to the aggregate loss rate, with a variation of 3%. This infers that UDP flow is fair and the loss rate is the same for the same load irrespective of the number of flows under the same packet size.

For comparison, Figure 6(b) shows the loss rate for zero and one hop with the same setting as the previous 4 hop settings. The same trend is found except with smaller loss rate. Both figures verify the results found in section IV-A.

V. MAC Layer Parameter Effects

In this section, we investigate the MAC layer parameters as potential sources of packet loss. This set of experiments use a fixed transmission power of 5dBm and MAC layer retransmissions are disabled for fixed transmission rates unless otherwise specified. There are four hops between source and destination in the experiment. In this section, we study the effects of different factors with a load of 8Mbps and using CBR packets of 1024 bytes at 1000 pkts/sec. Based on the measurements performed in our mesh testbed, we show that with multiple hops, packet loss can come from various points in the network. A change of the MAC layer parameter may reduce the packet loss at one place, but might lead to an increase at other places. Due to the dynamic wireless environment, the trend of the packet loss rate presents to the users is not as clear as that of one hop case with the tuning of the MAC layer parameters.

A. Transmission Modulation Rate

Figure 7(a) shows the loss rate by varying the transmission modulation rate under two different retransmission limits with the option of enabling/disabling RTS/CTS for a 4-hop network. The single hop result in Figure 7(b) shows the packet loss rate under the same setting as the 4-hop flows. There is a slight packet loss at the transmission rate 54Mbps. It reflects a trade off between sending the packet faster to relieve the effect of the limited buffer size and the higher probability of packet decode error to utilize high modulation rate. The

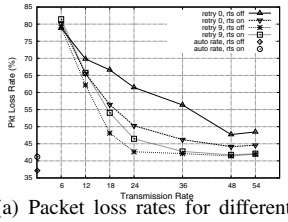

(a) Packet loss rates for different transmission modulation rates, retransmission limits and RTS/CTS settings.

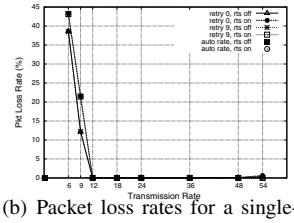

(b) Packet loss rates for a single-hop scenario for different transmission modulation rates, retransmission limits and RTS/CST settings.

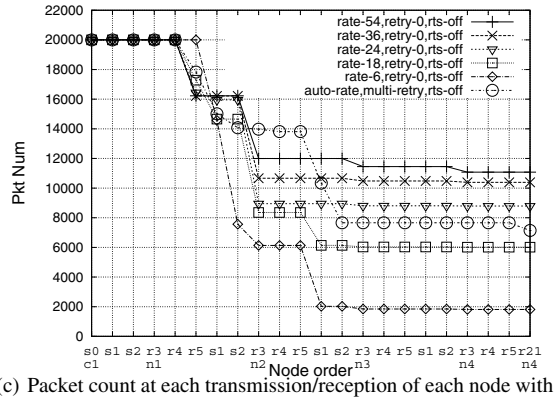

(c) Packet count at each transmission/reception of each node with RTS/CTS off.

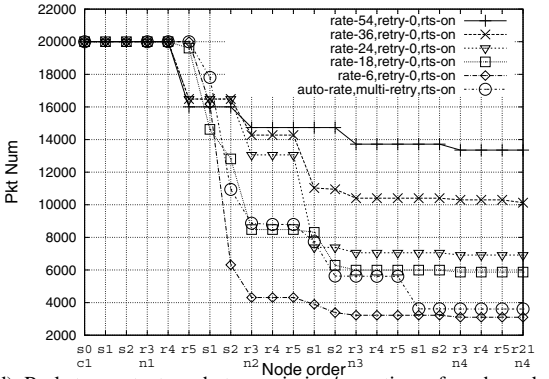

(d) Packet count at each transmission/reception of each node with RTS/CTS on.

Fig. 7. Experimental results for different transmission modulation rates.

packet loss rate becomes zero when the the transmission rate is greater than 12Mbps. We did a simple calculation to explain the reason. When sending l byte data packets, we assume the overhead incurred is given by:

- UDP/IP header 8+20 = 28 bytes
- MAC header + ACK = 38 bytes
- MAC/PHY procedure overheader = 400μs
  - DIFS (34μs), SIFS (16μs), preamble (20μs)
  - contention (approx. 310μs)

Given that the theoretical wireless capacity is C Mbps, the throughput of one frame is:

$$T(l, C) = \frac{8l}{400 \times 10^{-6} + (66 + l)\frac{8}{C \times 10^{-6}}} \qquad (1)$$

In order to sustain a packet generation rate of g pkts/sec, the capacity C of the wireless link must satisfy this inequality:

$$T(l, C) \geq 8lg \quad (2)$$

such that the throughput must be higher than the rate the frames are generated at.

For a generation rate of 1000 pkts/sec and a packet size of 1024 bytes, the capacity C should be greater than 14.53Mbps. Figure 7(b) shows that at the 12Mbps rate, packet losses are at a minimum. The 14.53Mbps value is higher because of the high average contention time assumption. The real experiment with 12Mbps does not have any interference in the one-hop setting. Rates lower than 12Mbps suffer from packet losses due to backlogging in the system. Rates higher than 12Mbps are fine since the capacity of the network is sufficient in sustaining the rate.

The Madwifi driver utilizes the SampleRate algorithm for rate control [13]. SampleRate sends most data packets at the bit rate that would provide the highest throughput based on estimates of the expected per-frame transmission time at each rate. If in the sampling process, no successful acknowledgment is received or the number of packets sent is a multiple of 10 on a specific link, it transmits packets with the highest rate which has not failed for 4 successive times.

Figure 7(c) and Figure 7(d) depicts the number of packets passing through the network at different modulation rates with RTS/CTS on or off, respectively.

As seen from the points between *n1:s1* and *n1:s2* at `node1`, and between *n2:s1* and *n2:s2* at `node2` in Figure 7(c), the 6Mbps rate and auto rate configuration shows packet drop in the transmission queue. The reason for the 6Mbps case is due to the low transmission speed to empty the queue slowly. For the auto rate case, the multi-retry mechanism increases the transmission time. However, according to the point between *n1:s2* and *n2:r3* at the second hop, rate 6Mbps and auto rate show much less losses in the medium than other rates. But there are more packets lost in putting into the transmission queue than can be saved with lower modulation rates. Even though there are some packet losses between *n2:s2* and *n3:r3* at the third hop, and between *n3:s2* and *n4:r3* at the fourth hop for 54Mbps flow, the number of packet loss from medium is still fewer than that of the drops in the transmission queue for the 6Mbps case. Combining the losses occurred somewhere else, higher rates is still a better choice for multi-hop flows with RTS/CTS disabled.

Similar trends can be seen in Figure 7(d). Moreover, with RTS/CTS enabled, the packet drops in the transmission queue deteriorate, as seen from the points between *n1:s1* and *n1:s2*, at at `node1` and between *n2:s1* and *n2:s2* at `node2`. We will have a comparison in Section V-C.

In general, we found that the packet loss rate decreases with higher transmission rate. The reason is that it decreases the air time for a packet by applying higher transmission rate, which helps mitigate the packet loss from limited buffer size or high load. As SampleRate is aggressive in utilizing high transmission rate, the performance is similar to that of middle-high range transmission 24-48Mbps, which is shown in Figure 7(a) and Figure 7(b).

B. MAC layer Retransmission Limit

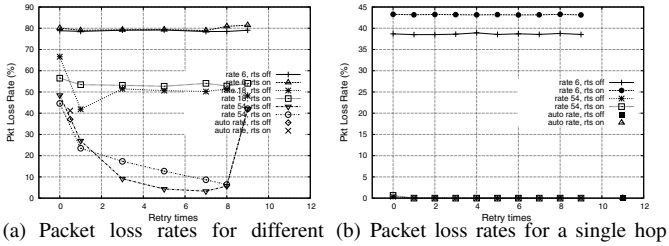

(a) Packet loss rates for different retransmission limits under different rates and RTS/CTS settings.

(b) Packet loss rates for a single hop scenario.

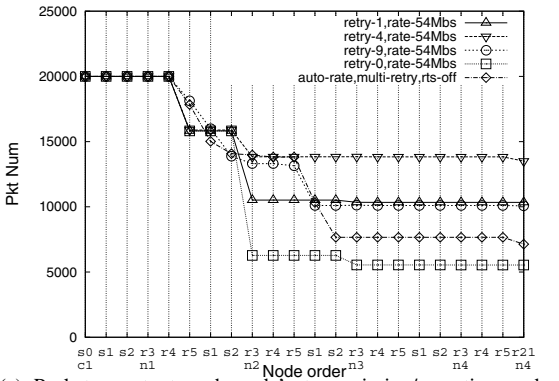

(c) Packet count at each node's transmission/reception under different retransmission limits.

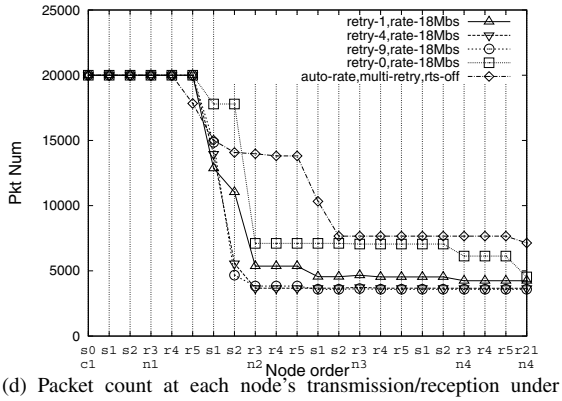

(d) Packet count at each node's transmission/reception under different retransmission limits.

Fig. 8. Experimental results for different MAC layer retransmission limits.

Figure 8(c) and Figure 8(d) depicts the number of packets passing through the network and the individual nodes for different maximum MAC retransmission limits at 54Mbps and 18Mbps rates respectively.

As seen from the points between *n1:s1* and *n1:s2* at `node1`, and between *n2:s1* and *n2:s2* at `node2` in Figure 8(c), having a maximum value of retransmission limit 9 at rate 54Mbps and automatic rate configuration show packet drops in the transmission queue while there are less drops for other retransmission limits. The reason for this is having higher retries took more time on the air which built up the transmission queue

faster. However, it was compensated by the less medium loss as seen from the points between *n1:s2* and *n2:r3* at the second hop. With zero or one retries at 54Mbps, the packet loss in the medium is considerably large. A small number of maximum retransmission limit of 4 presents a good choice for high rate 54Mbps, as the high transmission speed can allow for more retries before the transmission queue builds up while more retransmissions can save more packets from the medium loss.

With 18Mbps rate in Figure 8(d), as seen from the points between *n1:s1* and *n1:s2* at at `node1`, the amount of packet drops in the transmission queue increases with the increase with the retransmission limit. As SampleRate utilizes auto-rate multi-retry mechanism and prefer more retries on the high rate, there is smaller packet drop in the queue because its relatively high transmission speed can compensate for the increase of the transmission time due to more retries. High reliability at lower rates do not matter since the lower speed causes packet losses in the transmission queue. A higher retransmission limit won't help for low transmission rates either as seen from the points between *n1:s2* and *n2:r3* at the second hop.

Multi-hop flows that use high rate links will have better performance even links used a retransmission limit of around 4. In the case for low rate links, zero or one retransmission limit is better.

C. RTS/CTS

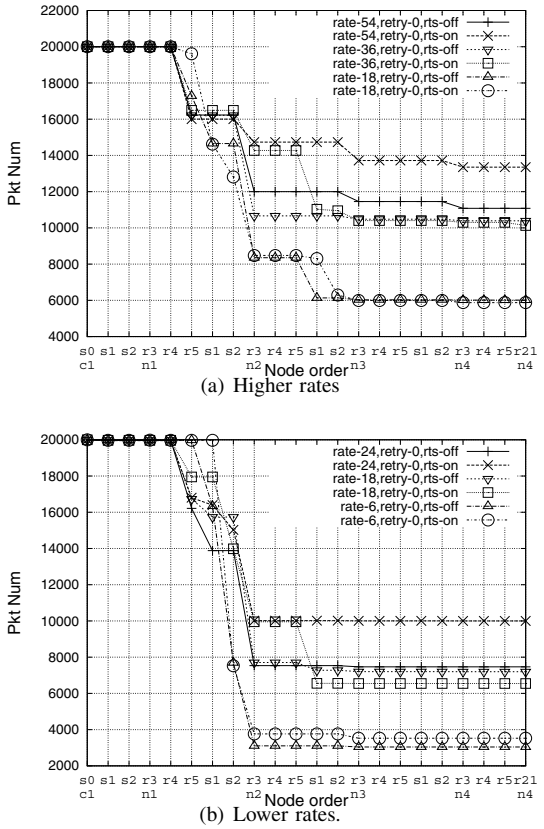

(a) Higher rates

(b) Lower rates.

Fig. 9. Packet count at each node's transmission/reception with and without RTS/CTS settings on.

Figure 9(a) and Figure 9(b) depicts the number of packets passing through the network and the nodes using the RTS/CTS mechanism with higher-range rates and lower-range rates respectively.

As seen from the points between *n1:s1* and *n1:s2* at `node1`, and between *n2:s1* and *n2:s2* at `node2` in Figure 9(a), turning on RTS/CTS for the 18Mbps rate configuration shows packet drop in the transmission queue while there are few drops for other rates. The reason is that high transmission rates greater than 18Mbps can empty the queue fast even with RTS/CTS mechanism on. According to the points between *n1:s2* and *n2:r3* at the second hop, even though with RTS/CTS protection, rate 18Mbps setting shows less medium loss than other rates, the combined effect to reach `node2` is almost the same as that with RTS/CTS disabled in which case there are more packets lost in the medium. However, without packet drop in the transmission queue, 36Mbps and 54Mbps settings can have more packets received from the air with the RTS/CTS protection according to point between *n1:s2* and *n2:r3* at the second hop.

For lower-range rates in Figure 9(b), as seen from the points between *n1:s1* and *n1:s2* at at `node1`, the 6Mbps rate setting drops a considerable amount of packets in the transmission queue due to its lowest transmission speed, which makes the effect of RTS/CTS mechanism not obvious. With 18Mbps and 24Mbps, turning on RTS/CTS makes more drops in the queue, however, it was compensated by the less medium loss as seen from the points between *n1:s2* and *n2:r3* at the second hop.

In summary, higher rates with RTS/CTS on under a small retransmission limit is a better choice for multi-hop flows and and the rates below 18Mbps won't see much difference as the large amount of packet drop in the transmission queue overwhelms the effect of turning on RTS/CTS mechanism. This is also reflected in Figure 7(a) and Figure 8(a).

D. Improving Reliability Through Manipulating System Parameters

After identifying the main sources of packet loss, in this section we outline the potential remedies to mitigate the packet loss due to the inappropriate system configuration. We evaluate the improvement by increasing the size of transmission queue (`txqueue`), the backlog queue (`backlog`) and the socket receive queue (`rmem`).

The setting for Figure 10(a) is a UDP CBR flow at 1000 pkts/sec with packet size 1024bytes each, under 5dbM transmission power, 6Mbps transmission rate, maximum 9 retries, and RTS/CTS off. As seen from the points between *n1:s1* and *n1:s2* at `node1`, increasing the transmission queue size can decrease the packet loss at the queue. However, with loss occurring at other places, the improvement is obvious for 2 and 3 hops, but not for 4 hops.

The setting for Figure 10(b) is a UDP CBR flow at 1500 pkts/sec with packet size 768bytes each, under 5dbM transmission power, 54Mbps transmission rate, no retries, and RTS/CTS off. From Figure 10(b), we can see the packets losses decrease by increasing the size of maximum backlog

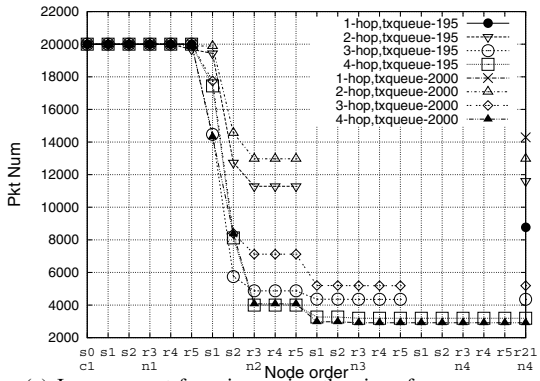

(a) Improvement from increasing the size of `txqueue`.

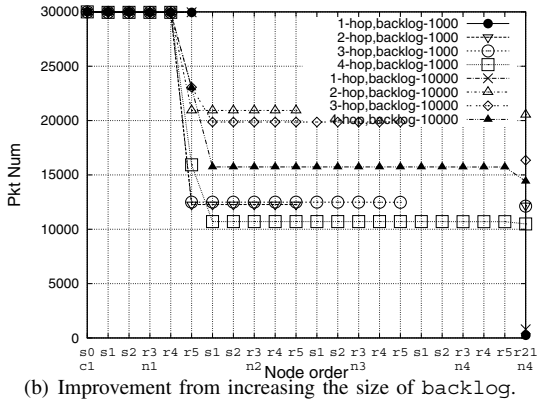

(b) Improvement from increasing the size of `backlog`.

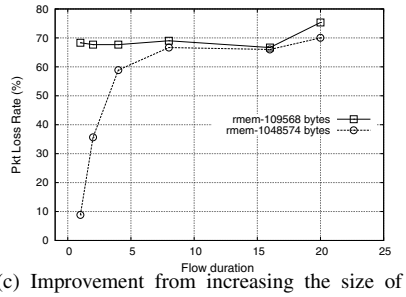

(c) Improvement from increasing the size of `rmem_max`.

Fig. 10. Experimental results for improvements through system parameters.

value, which is the number of unprocessed packets received from the link layer that the kernel can hold for upper layer processing, from 1000 to 10000. According to the difference between points *n1:r4* and *n1:r5* for multi-hop flows, the packet loss due to this drop is about 50% to 60% using the default maximum backlog value 1000, while the packet loss at this place is decreased to 20% to 30%.

We also run experiments to see the improvement from increasing the socket receiver buffer size from 109568 bytes to 1048574 bytes for a one-hop flow at 1500 pkts/sec with packet size 768 bytes each, under different flow duration. With such buffer increase, the packet loss rate for 1s flow can be decreased from 70% to 10%. As seen from Figure 10(c), even with a large buffer size, the loss rate increases considerably with time, as the slow application processing speed and frequently interrupted processing kept the socket receive queue full almost all the time after the first second.

VI. Conclusion

We performed a systematic study on the impacts of different factors on the unreliability of the mesh networks and the sources causing such unreliability. We focused on using the packet loss rate as the metric for reliability in this paper. The factors that we studied include traffic load, number of hops and flows, transmission rate, maximum retransmission times and RTS/CTS mechanism. In addition to performing the experiments on multiple hops, we include the results of one hop for comparison. In identifying the sources of packet loss, we developed FlowPaC tool to collect flow-based statistics from the nodes and analyzed the loss variability across time. Finally, we showed potential remedies to mitigate the high packet loss rate. System configurations have a great impact on packet losses when there is high contention for the computer resources from the heavy load, or a not-favored MAC layer configuration from low transmission speed, high retries or enabling RTS/CTS protection.

Acknowledgments

This research was supported in part by the National Science Foundation through the grants CNS-0831914 and CNS-0709264, and by the US Army Research Office through the MURI grant W911NF-07-1-0318.

References

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

[2] A. Sheth, S. Nedevschi, R. Patra, S. Surana, E. Brewer, and L. Subramanian, "Packet loss characterization in wifi-based long distance networks," in *InfoCom*, 2007.

[3] M. Rodrig, C. Reis, R. Mahajan, D. Wetherall, and J. Zahorjan, "Measurement-based characterization of 802.11 in a hotspot setting," in *E-WIND*, 2005.

[4] A. Scaglione and Y.-W. Hong, "Opportunistic large arrays transmission in wireless multihop ad hoc networks to reach far distances."

[5] J. Laneman and G. Wornell, "Exploiting distributed spatial diversity in wireless networks," in *Proc. Allerton Conf. Communications, Control, and Computing, (Monticello, IL)*, 2000. [Online]. Available: citeseer.ist.psu.edu/laneman00exploiting.html

[6] S. H. Y. Wong, S. Lu, H. Yang, and V. Bharghavan, "Robust rate adaptation for 802.11 wireless networks," in *MobiCom*, 2006.

[7] G. Holland, N. Vaidya, and P. Bahl, "A rate-adaptive MAC protocol for multi-Hop wireless networks," in *MobiCom*, 2001.

[8] A. K. Miu, H. Balakrishnan, and C. E. Koksal, "Improving Loss Resilience with Multi-Radio Diversity in Wireless Networks," in *MOBICOM*, 2005.

[9] A. Tsirigos and Z. Haas, "Analysis of multipath Routing-Part I: the effect on the packet delivery ratio," *Wireless Communications, IEEE Transactions on*, no. 1, 2004.

[10] Y. Yuan, H. Yang, S. Wong, S. Lu, and W. Arbaugh, "Romer: Resilient opportunistic mesh routing for wireless mesh networks," in *WiMesh*, 2005.

[11] J. Liu and S. Singh, "ATCP: TCP for mobile ad hoc networks," *IEEE J–SAC*, vol. 19, no. 7, pp. 1300–1315, 2001. [Online]. Available: citeseer.ist.psu.edu/liu99atcp.html

[12] "Soekris engineering," http://www.soekris.com/.

[13] "Madwifi:multiband atheros driver for wifi," http://www.madwifi.org/.

[14] "Thrulay: throughput and delay measurement tool," http://shlang.com/thrulay/.

[15] "TCPDUMP." [Online]. Available: http://tcpdump.org

[16] "D-itg, distributed internet traffic generator." [Online]. Available: http://www.grid.unina.it/software/ITG/