

The Impact of Link-Layer Retransmissions on Video Streaming in Wireless Mesh Networks (Invited Paper)

An Chan[†] Sung-Ju Lee[‡] Xiaolin Cheng[†] Sujata Banerjee[‡] Prasant Mohapatra[†]

[†]Department of Computer Science, University of California, Davis, CA 95616

[‡]Multimedia Communications & Networking Lab, Hewlett-Packard Labs, Palo Alto, CA 94304

[†]{anch,xlcheng,pmohapatra}@ucdavis.edu, [‡]{sjlee,sujata.banerjee}@hp.com

ABSTRACT

Link-layer retransmission is a feature of IEEE 802.11 standard protocol that aims to increase the reliability of data communications. However, when successive retransmissions fail, they add to the traffic congestion, raise the collision probability, and increase the end-to-end delay. Using our 4-hop wireless mesh network testbed, we evaluate the impact of link-layer retransmissions on the performance of video streaming in wireless multi-hop environment. Our experimental results show that when the traffic load is near or exceeds the network capacity, retransmissions cause erratic video quality and increase the end-to-end delay tremendously. When the best-effort traffic coexists, increasing the number of retransmissions degrades the goodput of best-effort traffic and increases the end-to-end delay of video streaming. Retransmissions add reliability and increase the video streaming quality only when the traffic volume is far below the network capacity limit.

Categories and Subject Descriptors

C.2.1 [COMPUTER-COMMUNICATION NETWORKS]: Network Architecture and Design—*Wireless communication*

General Terms

Experimentation

Keywords

Link-layer retransmissions, video streaming, multi-hop, wireless mesh networks

1. INTRODUCTION

With the wide deployment of wireless LANs (Local Area Networks), providing Quality of Service (QoS) in WLANs has been an active research topic. Many bandwidth-consuming and delay-sensitive applications, such as multimedia streaming and VoIP (Voice over IP), require network QoS to provide guaranteed bandwidth and bounded delay. Providing QoS in wireless networks, especially in multi-hop mesh networks is quite challenging and different from

wired networks [8]. The dynamic nature of wireless link quality and the broadcasting characteristic of the shared wireless medium increase the complexity of QoS provisioning. Link-layer retransmissions used in IEEE 802.11 MAC protocol [13] could be considered as a QoS feature. By retransmitting the frame that failed to be delivered, wireless errors can be overcome and data communication reliability can be improved. However, when retransmissions fail, they increase congestion, collision, and latency, and result in worsening the network performance. In this paper, we investigate the impact of link-layer retransmissions on the quality of video streaming over multi-hop wireless mesh networks.

Using our indoor wireless mesh network testbed, we evaluate the video streaming performance over multi-hop wireless networks in various network environments. Specifically, we measure the impact of IEEE 802.11 link-layer retransmissions on the video streaming quality when interference from intra-flow, inter-flow, and best-effort data traffic are present. We evaluate the video quality in PSNR (Peak Signal to Noise Ratio), PSNR variation, loss rate, latency, and delay jitter. Our experimental results show that when the network traffic volume is far below the network throughput capacity limit, retransmissions improve the video streaming quality. We observe that allowing simply *one* retransmission significantly improves the video quality over when no retransmission is used. However, increasing the retransmission limit beyond one or two does not generate significant performance gain. On the other hand, when the traffic load is near or beyond the network capacity limit, the effectiveness of retransmissions is minimal. In fact, in that scenario, increasing the retransmission limit results in a fluctuated streaming video quality and large end-to-end delay (more than five seconds on a 4-hop path with interference from a UDP flow).

The remainder of the paper is organized as follows. Section 2 introduces the related work. Section 3 describes our testbed and experimental setup. Experimental study of the impact of retransmission on video streaming under various scenarios is presented in Section 4. Section 5 concludes the paper.

2. RELATED WORK

There have been numerous studies that evaluate the effect of retransmission on wireless networks and multimedia over wireless multi-hop networks. The effect of retransmissions on multimedia transmission over WLANs was evaluated in [9]. Through simulations, they found that the effective throughput of wireless networks decreases as the retransmission limit increases. They developed a packet loss probability based analytical model to verify their findings. Our work differs in that we are focused on multi-hop networks and evaluate in a real testbed.

There have been efforts on utilizing different retransmission strategies in WLAN. Fast retransmission [11] and fragment-based re-

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

WICON'08 November 17-19, 2008, Maui, Hawaii, USA.

Copyright 2008 ACM ICST 978-963-9799-36-3 ...\$5.00.

transmission [17] were proposed to be used on top of IEEE 802.11e [14] networks to reduce latency. To achieve this, the fast retransmission scheme modifies the backoff procedure and sets the retransmission limit to three for VoIP traffic. The fragment-based scheme fragments the data frame to four fragments and appends CRC for each fragment. Network coding [16] was utilized to combine multiple lost packets to reduce the number of retransmissions in [21]. A heuristic for cooperative retransmission between the sender and neighbors that overhear the transmission was proposed [25] to decrease latency. All of the above proposals are designed for single-hop wireless networks, and it is unknown how effective these schemes will be in multi-hop mesh networks. For instance, as we will see later in our mesh testbed, we do not always see the incremental negative effect of increasing retransmissions on streaming video quality as observed in [9].

The impact of link-layer retransmissions on multi-hop wireless mesh networks has also been studied recently. It was found in [24] that the TCP performance is degraded drastically as the retransmission limit increases when the multi-hop network is heavily loaded. The effectiveness of retransmissions (no retransmissions or seven retry limit) on multimedia traffic was investigated on a wireless mesh network testbed [22]. It was found that although the retransmissions effectively reduce the packet loss rate, they vastly increase the end-to-end delay. An optimization framework for video streaming over multi-hop mesh network was proposed in [3] by considering the modulation rate (PHY layer), retransmission limit (MAC layer), routing (network layer), and packet scheduling (application layer). Using ETX (Expected Transmission Count) [10], a utility function was derived to calculate the retransmission limit based on estimated bandwidth and error rate. This scheme however has large messaging overhead and system complexity issues, as noted by the authors.

We focus our attention on the effect of link-layer retransmissions on multimedia transmissions, in particular streaming video, over an indoor multi-hop wireless mesh network testbed. We create various network scenarios by varying the path length and video coding rate and evaluate the retransmissions in the presence of intra-flow interference. We also study the effectiveness of retransmissions with inter-flow interference and when the streaming video coexists with the best effort traffic. In addition to the network statistics such as end-to-end delay and packet loss rate, we also assess the video quality using Peak Signal-to-Noise Ratios (PSNR), which is a widely used metric in the media community.

3. WIRELESS MESH NETWORK TESTBED

3.1 Testbed Topology

Our testbed consists of mesh routers and mesh clients [2]. A mesh client establishes wireless links using mesh routers as relays for data communication. Mesh routers establish wireless links among themselves to provide connectivity for the entire wireless mesh network. Mesh routers essentially have two functions; (i) to provide interfaces for the mesh clients to connect to the mesh network, and (ii) to run a routing protocol to forward packets toward destinations in the mesh network. As our study focuses on the impact of link-layer retransmissions, we do not implement a routing protocol in our testbed. Instead, we assume static routing and use WDS (Wireless Distributed System) for packet forwarding. Figure 1 shows typical 2-hop, 3-hop, and 4-hop networks built from our testbed. The dotted lines represent wireless links. In our experiments, a video stream is sent from the one end of the network to the other end. In Figure 1, M1, M2 and M3 are three mesh routers while C1 and C2 are mesh clients.

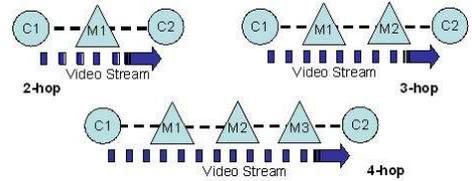


Figure 1: Our wireless mesh network testbed topology.

3.2 Experimental Setup

In our testbed, mesh routers are Soekris board APs, and mesh clients are HP compaq nc6000 laptops equipped with HP W500 802.11a/b/g wireless LAN cards (using Atheros chipsets). The Soekris boards and HP laptops are installed with Linux operating system with kernel version 2.6.22.1. MadWifi [18] is installed in the nodes as WLAN drivers. We modify the MadWifi source to change the retransmission limit. We configure the wireless mesh network with IEEE 802.11a [12] to reduce interference from co-existing 802.11b/g networks in the building. All nodes, both mesh routers and mesh clients, in our network operate in the same channel, Channel 36 (5.18GHz). Moreover, all nodes are in the same contention region. As all nodes can “hear” each other, there are no hidden nodes [15, 23]. We set the modulation rate of the data transmission to 6 Mbps. We disable rate adaptation as it will vary the network condition and complicate analyzing the results.

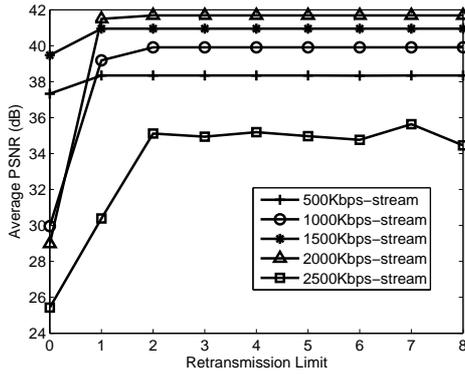
We use a 2000-frame *highway* clip for video streaming. The video clip is coded into MPEG4 streams using *ffmpeg* [4] with a frame rate of 25fps (i.e., the clip lasts for 80 seconds). UDP/RTP is used for the streaming protocol. We evaluate the quality of video streaming by calculating the Peak Signal-to-Noise Ratio (PSNR), which is the most commonly used metric for video quality. At the receiver, we compare the received video clip with the original clip and calculate the PSNR (i.e., we use a Full Reference Method). Note that the PSNR calculation is based on the mean squared error (MSE) of two images in corresponding video frames from the original clip and the received (distorted) clip. It does not take into account the end-to-end delay or delay jitter. In other words, PSNR in our experiments directly reflects the objective video quality when an infinite playback buffer at the receiver is assumed. Nevertheless, PSNR is still an important metric for evaluating video streaming quality. We also measure other network metrics such as end-to-end latency and packet loss rate of the streamed video to comprehend the effectiveness of retransmissions.

4. EXPERIMENTAL RESULTS

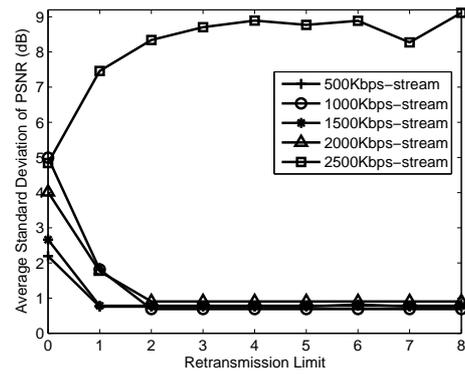
4.1 Intra-Flow Interference

We first send only one video stream from C1 to C2 with different path lengths (see Figure 1). In multi-hop transmissions, when all the wireless links in a path operate on the same channel, the transmission in one link could interfere with the transmissions in other links in the path. This type of interference is referred to as “intra-flow interference.”

In this experiment, we stream the video over 2-hop, 3-hop and 4-hop wireless paths in our testbed. For each n -hop path (where $n = 2, 3, 4$), we use five different video coding rates of 500, 1000, 1500, 2000, and 2500 Kbps. The source (C1) sends the video in a transmission rate near the coding rate. As the coding rate increases, the transmission rate will approach or exceed the throughput capacity of a particular n -hop path. The higher coding rate incurs the higher intra-flow interference level. We apply different retrans-

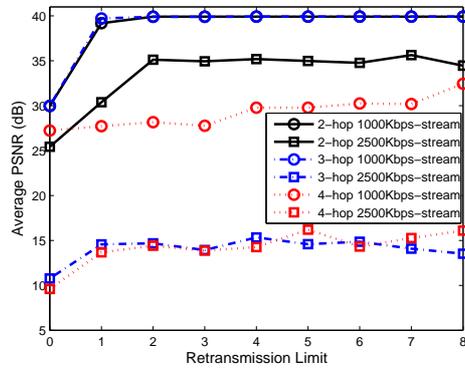


(a) PSNR.

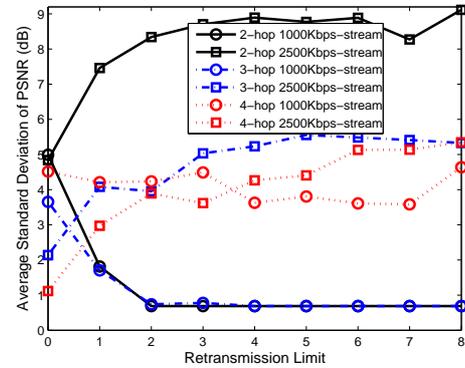


(b) Standard deviation of PSNR.

Figure 2: Video quality with varying coding rate and retransmission limit on a 2-hop path.



(a) PSNR.



(b) Standard deviation of PSNR.

Figure 3: Video quality with varying coding rate, retransmission limit, and path length.

mission limits for each coding rate to study the effectiveness of retransmissions in various levels of interference.

4.1.1 PSNR Video Quality with Varying Rates

We first evaluate how retransmissions improve the video performance with various coding rates, without varying the path length. Figure 2(a) shows the average PSNR of the video streaming over a 2-hop path when retransmission limits and video coding rates are varied. We can observe the effectiveness of retransmissions, as the PSNR of the video streaming is the lowest when the retransmission limit is zero, for all coding rates. As the end-to-end throughput capacity of our 2-hop path is near 2500 Kbps, we can categorize the coding rates into two. For the coding rates from 500 Kbps to 2000 Kbps, which are within the capacity of the path, their PSNR increases with the increasing coding rate. This is an expected result since higher coding rates compress the original clip less and hence show smaller distortion and higher PSNR.

We can see that the PSNR improves with retransmissions, and stabilizes after a few retransmission limits of one or two. Note that in our testbed, all nodes are within the contention region of each

other and there are no hidden nodes. Therefore, the packet loss or distortion in videos of the coding rates from 500 Kbps to 2000 Kbps is due to the background noise or packet collisions between the nodes counting down backoff timer to zero at the same time. The probability of the latter is very small in a network with three nodes. Hence, only a very small number of the video frames are lost or distorted, and they can be recovered by few retransmissions.

For the 2500 Kbps stream, the transmission rate has reached the throughput capacity limit of a 2-hop path. Moreover, when the transmission rate is high, the intra-flow interference level is also high. Queuing delay of the packets increases and consequently, there are buffer overflow which results in dropped video frames. Increasing the retransmission limit could help recover certain video frames and improve PSNR. However, it also worsens the traffic congestion and causes some video frames to have very low PSNR. We can see from Figure 2(b) that the standard deviation of PSNR for the 2500 Kbps stream increases as the retransmission limit increases. This indicates that the quality of 2500 Kbps stream becomes erratic when more retransmissions are allowed. To summarize, with strong intra-flow interference, the increase in retransmis-

sion does not improve the quality of 2500 Kbps video, although the average PSNR increases. On the other hand, for video streams with lower coding rates, the standard deviation values quickly decrease and maintain the low values after allowing two or more retransmissions.

4.1.2 PSNR Video Quality with Varying Hops

We have repeated experiments over 3-hop and 4-hop paths and observed similar trends. Although we have results for five different coding rates, we only present 1000 Kbps and 2500 Kbps streams for readability in Figure 3(a).

The performance of the video streams largely depends on whether the throughput capacity limit of the path has been reached. As the number of hops increases, the capacity of the path decreases. It is more so in our testbed where all nodes are in the contention region of each other to compete for the wireless medium access. Using *iperf* [19], we found that the capacity of 3-hop and 4-hop paths are around 1500 Kbps and 1050 Kbps, respectively. The transmission rate of 1000 Kbps-streams in 2-hop and 3-hop paths is below the path capacity, and hence the video streams reach a very high and stable PSNR level (almost 40 dB) when the retransmission limit is set to one or two. We also find from Figure 3(b) that the PSNR standard deviations of the both 1000 Kbps streams quickly decrease to their minimum as the retransmission limit increases to two. In contrast, the 2500 Kbps streams in 3-hop and 4-hop networks exceed the network capacity limit. Many packets are lost due to buffer overflow and increasing the retransmission limit can only recover a small number of packets. The PSNR of the 2500 Kbps streams is always much inferior to the 1000 Kbps streams in 3-hop and 4-hop networks. The standard deviations also keep increasing as the retransmission limit increases. Hence, when the video exceeds the capacity and overwhelms the network, retransmissions have very little merit.

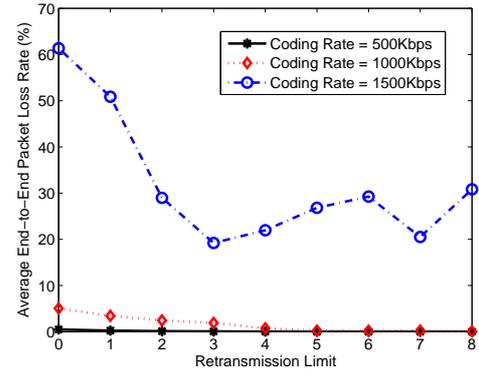
An interesting observation can be made for the 1000 Kbps stream in the 4-hop network. The transmission rate is close to, but does not exceed the network capacity of around 1050 Kbps. When we increase the retransmission limit, more retransmitted packets are in the network and after a certain point, the throughput capacity is exceeded. Hence, increasing the retransmissions does not recover the lost packets and the PSNR only slowly increases. The standard deviation also stays high when the retransmission limit increases. This result is similar to the 2500 Kbps streaming in the 2-hop network. In both cases, the transmission rates are very close to the network throughput capacity. As they do not exceed the capacity limit as much as 2500 Kbps streams in 3-hop and 4-hop networks do, their PSNR values are much superior than 2500 Kbps streams in 3-hop and 4-hop paths. However, increasing the retransmissions does not improve the video quality when the video coding rate is close to or exceeds the network throughput capacity.

For the rest of this paper, we focus the discussion and the experiments on the 4-hop network.

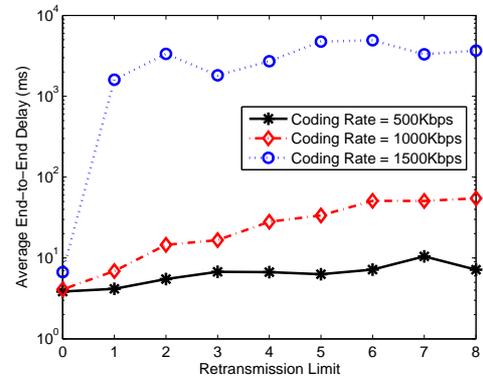
4.1.3 Other Performance Metrics

PSNR has been widely used for evaluating the video quality. However, it is not a perfect metric; it does not take into account end-to-end delay, and the average PSNR could still be high even with high standard deviation of PSNR. We hence assess the impact of retransmissions on video streaming with intra-flow interference in the 4-hop network, using popular metrics in the networking community: packet loss rate, end-to-end delay, and delay jitter.

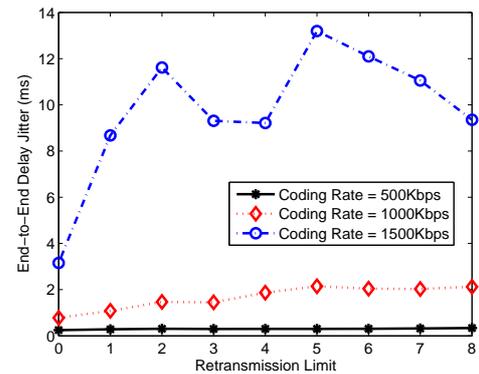
Figure 4(a) shows the average packet loss rate of three video coding rates as the retransmission limit increases. We see that for the 1500 Kbps stream, although the average packet loss rate generally



(a) Packet loss rate.



(b) End-to-end delay.



(c) End-to-end delay jitter.

Figure 4: Other performance metrics of video streams with intra-flow interference.

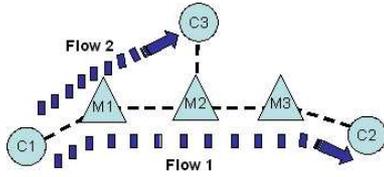


Figure 5: Testbed topology with two video flows.

decreases as more retransmissions are allowed, the overall packet loss rate is still very high and fluctuating as the retransmission limit further increases. For the coding rates of 500 and 1000 Kbps, as they are under the network capacity, the loss rate is quite low. For the 1000 Kbps stream, the loss rate slowly decreases with the increase in the retransmission limit, but is always still higher than that of the 500 Kbps stream. As discussed in Section 4.1.2, since the 1000 Kbps stream is near the network capacity, the increase in retransmissions causes the traffic to exceed the throughput capacity and some lost packet cannot be recovered. The PSNR of 1000 Kbps was always lower than that of the 500 Kbps stream (although 500 Kbps was not presented in Figure 3(a) for clarity).

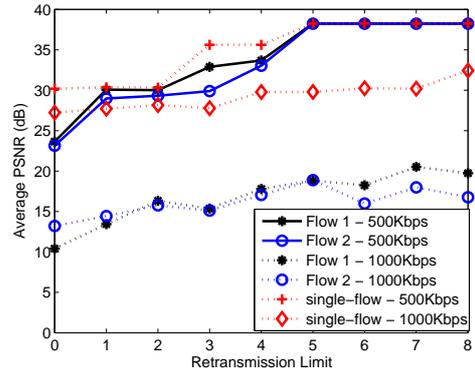
Figure 4(b) presents the latency of the video streams in logscale with different coding rates. The number of retransmissions has little or no effect on the 500 Kbps stream, as its loss rate is very low, and the coding rate is well below the network capacity. We see a slow increase for the 1000 Kbps stream as the retransmission limit increases. As the latency is measured only for packets that reach the destination, the packets recovered after retransmissions have increased delay. For the 1500 Kbps stream, the latency sharply increases when the retransmission limit is changed from zero to one. Increasing the retransmission adds to the network traffic and contributes to congestion, especially when the network already has high traffic load. This in turn will cause buffer overflow and increase the end-to-end delay. We will further study the relationship of latency and buffer overflow in Section 4.3 when the best-effort data traffic coexists with the video stream.

Figure 4(c) shows the average end-to-end delay jitter of the video streams. We observe similar trends as with the latency in Figure 4(b).

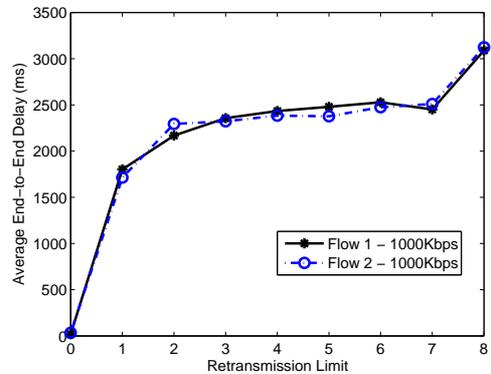
4.2 Inter-Flow Interference

In order to create inter-flow interference, we add one more client, C3 as a destination of the second flow to the 4-hop network as shown in Figure 5. In addition to a 4-hop video stream from C1 to C2 (Flow 1), we now have a 3-hop video stream from C1 to C3 (Flow 2). For this experiment, we start Flow 2 ten seconds later than Flow 1. Both flows stream video of the same coding rate and all nodes have the same retransmission limit. We use the video coding rates of 500 Kbps and 1000 Kbps so as not to overwhelm the network. We first evaluate in Figure 6(a) the PSNR of the two flows when different coding rates and retransmission limits are used.

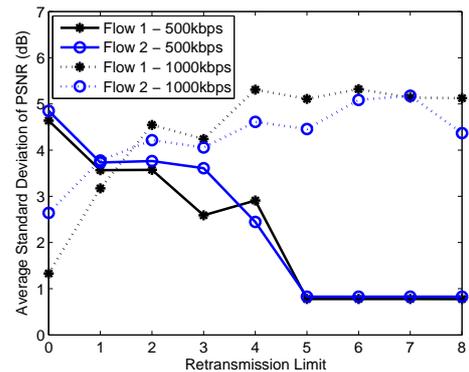
Similar to the intra-flow interference experiments, increasing the retransmission limit improves the PSNR of both flows in either coding rates. With 500 Kbps video streaming, both flows eventually reach the PSNR of around 38 dB, which was the highest a single 500 Kbps stream attained. When no retransmission is used, the PSNR of both 500 Kbps streams are lower (around 24dB) than the single-flow case (around 30dB). With two flows, there are inter-flow interference as well as intra-flow interference, and hence more frames are lost or distorted. The PSNR increases when the retransmission limit increases from zero to one, and matches that of the single flow scenario.



(a) PSNR.



(b) End-to-end delay.



(c) Standard deviation of PSNR.

Figure 6: Video quality with inter-flow interference.

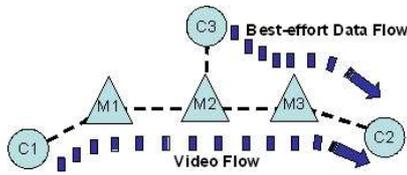


Figure 7: Testbed topology with video and best-effort traffic flows.

With two video streams of 1000 Kbps coding rate, the traffic volume exceeds the throughput capacity limit of a 4-hop network. Although the average PSNR of the both 1000 Kbps streams generally improves with the increase in the retransmission limit, we can see the big jump in latency in Figure 6(b) when the retransmission is turned on. Increasing the retransmission limit on a high traffic network could further increase the congestion in the network. This in turn increases the queuing time of a packet and hence the end-to-end delay. Such a long delay would be intolerable for real-time applications such as video conferencing. Furthermore, as Figure 6(c) indicates, the average standard deviation of PSNR for the 1000 Kbps videos increases and thus the video quality becomes unstable as the retransmission limit increases. Hence, the effectiveness of retransmissions is also questionable in a high traffic network with inter-flow interference.

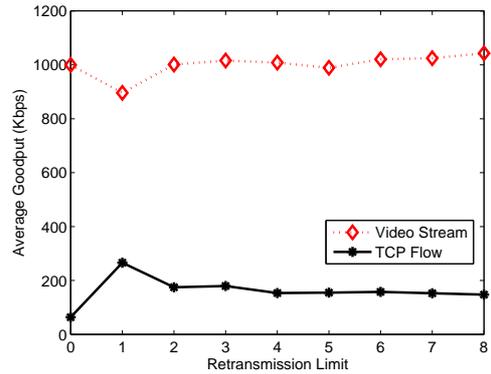
4.3 Coexistence with Best-Effort Traffic

It is common for video streams and best-effort data traffic (TCP or UDP) to coexist in the wireless network. We study the interaction of best-effort traffic and video streaming and the effectiveness of retransmissions in this scenario. The best-effort data traffic is sent from C3 to C2 in our testbed shown in Figure 7. The video flow starts first and the best-effort flow starts 10 seconds later, and lasts for 60 seconds. We use *iperf* [19] to generate best-effort traffic. For the video flow, we use a 1000Kbps-stream.

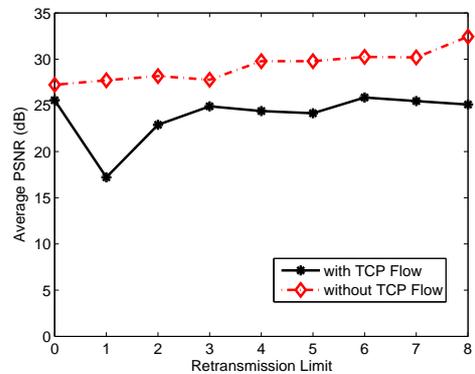
4.3.1 Interaction with TCP Traffic

We look at the impact of TCP traffic on video streaming when different retransmission limits are used. Figure 8(a) shows the goodput of the flows, while Figure 8(b) shows the PSNR of the video flow with and without the coexisting TCP flow. When no retransmissions are allowed to recover lost packets, TCP experiences high loss rate and decreases its transmission rate. The video flow, which uses UDP, takes advantage of this to grab a large portion of the network capacity. With more UDP traffic occupying the network, TCP's congestion control mechanism further decreases the transmission rate and the TCP session sometimes cannot even be established. The average time interval between two consecutive TCP packet transmissions is near 130 ms while it is mere 10 ms when a single TCP flow is the only traffic in a 3-hop path. The corresponding PSNR of the video when the retransmission limit is zero is very high, almost as high as when there is no TCP flow.

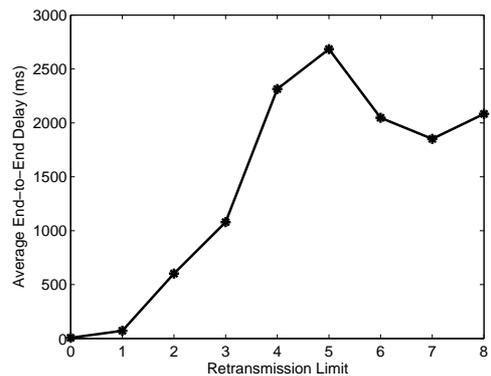
However, this unfairness alleviates when the link-layer retransmission limit increases to one. With the link-layer retransmissions recovering lost packets, TCP does not decrease the transmission rate as fast and obtains a larger part of the network capacity. We can see the increase of TCP goodput when the retransmission limit increases from zero to one in Figure 8(a). However, further increasing the retransmission limit does not increase the TCP performance, and in fact, decreases the goodput. Although some packets reach the destination after a large number of retransmissions, the ACK timeout has already expired at the source and the source



(a) Goodput.



(b) Video PSNR.



(c) End-to-end video delay.

Figure 8: Video and TCP performance when they coexist.

performs the transport layer retransmission. Moreover, multiple retransmissions cause the network to be congested and TCP congestion control will decrease the transmission rate. As TCP decreases the transmission rate, the video flow again occupies increasingly larger portion of the network capacity. We can also see from Figure 8(b) that the PSNR of the video stream increases when the retransmission limit is increased from one when TCP coexists.

Although the average PSNR of the video flow improves as the retransmission limit increases, the end-to-end delay sharply increases, as shown in Figure 8(c). The delay is already 600 ms when the retransmission limit is two. Such a long latency is unacceptable for delay-sensitive applications. We conjecture that the huge increase in end-to-end delay is caused by the queuing delay. When the retransmission limit increases beyond a certain value (five in this scenario), buffer overflow occurs in the intermediate nodes (i.e., M1, M2 and M3) as many retransmitted packets have been queued up. This results in many packet losses. The lost packets do not contribute to the calculation of end-to-end delay. Therefore, the end-to-end delay drops after the retransmission limit of five.

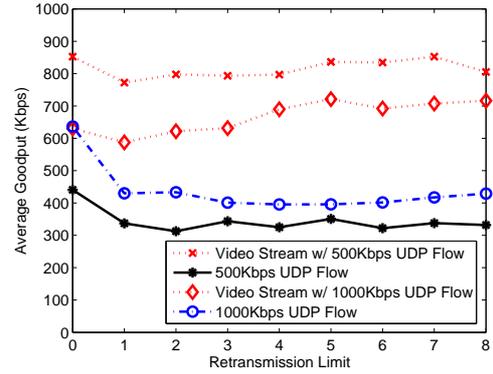
4.3.2 Interaction with UDP Traffic

When UDP traffic coexists with the video flow, the video flow can no longer dominate the medium access as it has with the TCP flow. As shown in Figure 9(a), when the retransmission limit is zero, the goodput of 1000 Kbps UDP flow and the video flow (1000 Kbps coding rate) are almost the same. When coexisted with the 500 Kbps UDP flow, as the video flow is more aggressive in generating traffic, it occupies a larger portion of the network capacity. As the retransmission limit increases, the video flow accesses more network capacity than the UDP flow, especially with the 1000 Kbps UDP flow. We can see the gap between the video stream and the 1000 Kbps UDP flow increases as the retransmission limit grows.

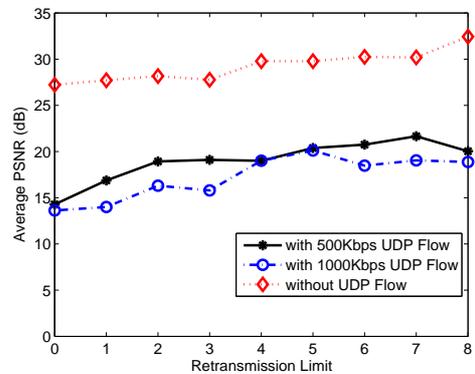
From the network topology shown in Figure 7, we know that the last two hops of the best-effort data flow and the video flow share the same links (M2-M3-C2). All UDP data packets to M2 are from C3, while the video frames are from C1 to M1, and then from M1 to M2. As the three links, C1 to M1, M1 to M2 and C3 to M2, are all in the same contention region, they have to compete with each other for the network capacity. Video flow occupies the network capacity first because we do not start the UDP flow until 10 seconds after starting the video flow. Once the UDP flow starts, if the three links always have packets to send (although packets from M1 depends on C1, M1 already has a lot of packets to send since we have already ran the video stream for 10 seconds), they will equally divide the network capacity. When retransmission is allowed, a lot of video packets will be generated from C1 to M1 and from M1 to M2. The extra video packets induced from these two links are more than the UDP packets induced from the link C3 to M2. Probabilistically, more packets from C1 (video frames) can get to M2 compared with packets from C3. So, the goodput of video stream is larger than that of the UDP flow. This goodput imbalance becomes more obvious as the retransmission limit increases.

We observe in Figure 9(b) that the PSNR of the video flow drops significantly when there is interference from the UDP flows. Increasing the retransmission count improves the video quality, but is nowhere near when there is no UDP traffic, and is also inferior to when the video coexists with a TCP traffic.

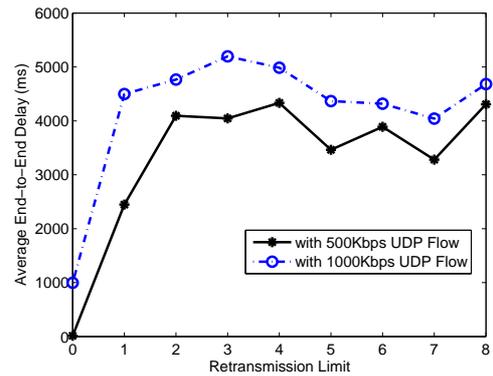
Figure 9(c) shows that increasing the retransmission limit quickly enlarges the end-to-end latency of the video flow. It shows a similar trend as when a TCP traffic coexisted. The video latency when competing with the UDP flow is even larger as the UDP flow sends traffic much more aggressively than the TCP flow.



(a) Goodput.



(b) Video PSNR.



(c) End-to-end video delay.

Figure 9: Video and UDP performance when they coexist.

4.4 Summary and Discussion

Using our testbed experiments, we have studied the impact of the IEEE 802.11 link-layer retransmission on video streaming over various multi-hop wireless mesh network scenarios.

When only single video streaming flow exists in network and the video suffers only from intra-flow interference, retransmissions increase the video quality. In fact, most lost packets are recovered by one or two retransmissions. However, when the coding rate of the video flow exceeds the network capacity, increasing the retransmission limit will degrade the performance. Although the PSNR value might not drop, the standard deviation of PSNR, packet loss rate, latency, and jitter all increase as the retransmission limit grows.

When two video streams create inter-flow interference to each other, the video quality gradually improves as the retransmission limit grows. Similar to the single video flow case, if the combined transmission rates exceed the network capacity, increasing the retransmissions causes the video quality to be erratic and incurs large end-to-end delay.

When the video stream coexisted with TCP traffic, the video performance dropped when the retransmission limit was increased to one. The video PSNR and goodput improve when the retransmissions were allowed more than once. However, as the retransmission limit increases, there was a sharp increase in the latency. Hence, when coexisting with TCP traffic, for best video performance, turning off the retransmissions would be a smart choice, although at the expense of poor TCP performance. When UDP traffic coexisted with the video flow, the video PSNR slowly improved with the increasing retransmission limit. However, larger number of retransmissions do not improve the goodput of the video streaming or the UDP flows. As retransmissions caused a quick increase in latency, the effectiveness of retransmissions when the video stream coexists with the UDP flows is questionable.

With wider deployment of wireless mesh networks for both research and practice [5, 1, 20], the users are not satisfied with just wireless connectivity, and also demand QoS requirement to be met. Retransmissions is one way to provide QoS in wireless networks. Our study showed that when the network is overwhelmed with traffic, retransmissions lose effectiveness and only add to the congestion. Therefore, admission control [7, 26, 6] should be implemented and combined with retransmissions to improve QoS.

Our testbed experimental study also suggests simply fine-tuning the retransmission limit parameter can greatly improve the network performance. Based on our observations, we should consider using different retransmission limits for the video flow and the best-effort flow. Other future directions include adaptive retransmission where packets are retransmitted based on the network condition. For instance, no retransmissions should be allowed when the network is congested. Selective retransmission is another scheme to investigate, where packets are retransmitted based on the type of video frames. As surveyed in Section 2, other retransmission strategies have also been proposed. It is important for the research community to implement the new schemes in real wireless network systems for thorough assessment and wide deployment.

5. CONCLUSION

Through our experimental studies, we found that the impact of link-layer retransmission on video streaming over wireless mesh network heavily depends on the total traffic load in the network. If the network capacity limit is exceeded by the traffic rate, increasing the number of retransmissions only causes instability of the video quality and tremendously increases the end-to-end delay. When the best-effort data traffic coexisted with the video streaming, in-

creasing the retransmission limit also degrades the goodput of the best-effort traffic, especially TCP. Enabling more retransmissions improves the video quality when the traffic volume is far below the network capacity. But this advantage of retransmission could not be easily realized in today's heavily used wireless networks. With increasing demand of real-time multimedia communication on wireless mesh networks, we believe greater effort is needed to improve the QoS over wireless mesh networks.

6. REFERENCES

- [1] M. Afanasyev, T. Chen, G. M. Voelker, and A. C. Snoeren. Analysis of a Mixed-Use Urban WiFi Network: When Metropolitan becomes Neapolitan. In *Proceedings of IMC*, 2008.
- [2] I. F. Akyildiz, X. Wang, and W. Wang. Wireless mesh networks: a survey. *Computer Networks*, 47(4):445–487, March 2005.
- [3] Y. Andreopoulos, N. Mastrorade, and M. V. D. Schaar. Cross-layer optimized video streaming over wireless multihop mesh networks. *IEEE Journal on Selected Areas in Communications*, 24(11):2104–2115, November 2006.
- [4] F. Bellard. FFMPEG(<http://ffmpeg.mplayerhq.hu/>), 2008.
- [5] J. Bicket, D. Aguayo, S. Biswas, and R. Morris. Architecture and Evaluation of an Unplanned 802.11b Mesh Network. In *Proceedings of MobiCom*, 2005.
- [6] I. D. Chakeres and E. M. Belding-Royer. PAC: Perceptive Admission Control for Mobile Wireless Networks. In *Proceedings of QSHINE*, 2004.
- [7] X. Cheng, P. Mohapatra, S.-J. Lee, and S. Banerjee. MARIA: Interference-Aware Admission Control and QoS Routing in Wireless Mesh Networks. In *Proceedings of ICC*, May 2008.
- [8] X. Cheng, P. Mohapatra, S.-J. Lee, and S. Banerjee. Performance Evaluation of Video Streaming in Multihop Wireless Mesh Networks. In *Proceedings of NOSSDAV*, June 2008.
- [9] S. Choudhury, I. Sheriff, J. D. Gibson, and E. M. Belding-Royer. Effect of payload length variation and retransmissions on multimedia in 802.11a WLANs. In *Proceedings of IWCMC*, pages 377–382, New York, NY, USA, 2006. ACM.
- [10] D. S. J. D. Couto, D. Aguayo, J. Bicket, and R. Morris. A High-Throughput Path Metric for Multi-Hop Wireless Routing. In *Proceedings of MobiCom*, 2003.
- [11] G.-H. Hwang and D.-H. Cho. Fast Retransmission Mechanism for VoIP in IEEE 802.11e Wireless LANs. In *Proceedings of VTC*, September 2004.
- [12] IEEE. IEEE 802.11a, Part 11: Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) specifications: High-speed Physical Layer in the 5 GHz Band.
- [13] IEEE. IEEE 802.11-1999, Part 11: Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) specifications, IEEE Std., August 1999.
- [14] IEEE. IEEE 802.11e/D8.0, Draft amendment to 802.11: Medium Access Control (MAC) Quality of Service (QoS) Enhancement, February 2004.
- [15] L. Jiang and S. C. Liew. Hidden-Nodes Removal and Its Application in Cellular WiFi Networks. *IEEE Transactions on Vehicular Technology*, 56(5):2641–2654, September 2007.
- [16] S. Katti, H. Rahul, W. Hu, D. Katabi, M. Medard, and

- J. Crowcroft. XORs in the Air: Practical Wireless Network Coding. In *Proceedings of SIGCOMM*, 2006.
- [17] C.-F. Kuo, H.-W. Tseng, and A.-C. Pang. A Fragment-Based Retransmission Scheme with QoS Considerations for Wireless Networks. In *Proceedings of IWCMC*, August 2007.
- [18] MadWifi. (Online – <http://madwifi.org/>), 2008.
- [19] NLANR/DAST. Iperf - The TCP/UDP bandwidth measurement tool (<http://dast.nlanr.net/Projects/Iperf/>), 2008.
- [20] J. Robinson, R. Swaminathan, and E. Knightly. Assessment of Urban-Scale Wireless Networks, with a Small Number of Measurements. In *Proceedings of MobiCom*, 2008.
- [21] E. Rozner, A. P. Iyer, Y. Mehta, L. Qiu, and M. Jafry. ER: Efficient Retransmission Schemem for Wireless LANs. In *Proceedings of CoNEXT*, December 2007.
- [22] Y. Sun, I. Sheriff, E. M. Belding-Royer, and K. C. Almeroth. An Experimental Study of Multimedia Traffic Performance in Mesh Networks. In *Proceedings of WitMeMo*, 2005.
- [23] F. Tobagi and L. Kleinrock. Packet Switching in Radio Channels: Part II – the Hidden Terminal Problem in Carrier Sensing Multiple Access and the Busy Tone Solution. *IEEE Transactions on Communications*, 23(12):1417–1433, December 1975.
- [24] Y. Wang, K. Yu, Y. Liu, and H. Zhang. Effects of MAC Retransmission on TCP erformance in IEEE 802.11-based Ad-hoc Networks. In *Proceedings of VTC*, May 2008.
- [25] L. Xiong, L. Libman, and G. Mao. Optimal Strategies for Cooperative MAC-Layer Retrasnmission in Wireless Networks. In *Proceedings of WCNC*, March 2008.
- [26] Y. Yang and R. Kravets. Contention-aware admission control for ad hoc networks. *IEEE Transactions on Mobile Computing*, 4(4):363–377, 2005.