Performance Analysis of the Intertwined Effects between Network Layers for 802.11g Transmissions

Jon Gretarsson, Feng Li, Mingzhe Li, Ashish Samant, Huahui Wu, Mark Claypool and Robert Kinicki
Dept. of Computer Science, Worcester Polytechnic Institute
Worcester, MA, 01609, USA
{jontg,lif,lmz,asamant,flashine,claypool,rek}@cs.wpi.edu

ABSTRACT

While the canonical behavior of today's home Internet users involves several residents concurrently executing diverse Internet applications, the most common home configuration is a single external connection into a wireless access point (AP) that promises to provide concurrent high-bandwidth Internet access for multiple hosts through a wireless local area network (WLAN). Recent research has attempted to assess the performance impact of hosts with weak wireless connectivity upon the other WLAN hosts by employing measurement studies or analytic models that focus primarily on wireless channel characteristics. This paper examines the intertwined effects on performance of the user applications, the network protocol and the wireless channel characteristics via carefully designed experiments that leverage previously developed network measurement tools. The study provides empirical evidence that suggests the overall performance of a WLAN is not only determined by the individual wireless channel qualities associated with each host, but also by the interaction of the various network layers with respect to transmission contention, queuing at the access point, transport protocol, and behavior of the specific applications. These results imply that effective WLAN performance modeling needs to include details on multiple network layers.

Categories and Subject Descriptors: C.2.m [Computer-Communication Networks]: Miscellaneous

General Terms: Measurement, Performance, Design.

Keywords: Streaming Media, Wireless, IEEE 802.11.

1. INTRODUCTION

It is increasingly common for home users to access the Internet via a wireless Local Area Network (WLAN) connected to a single external broadband or ADSL connection. Such a configuration allows several home residents to concurrently access the Internet while using a variety of Internet applications. Furthermore, growth in WLAN deployment at universities [9] increases the likelihood that concurrent wireless hosts will access the Internet through a common wireless access point (AP).

This growth in WLAN use has encouraged research in modeling [4, 5] and measuring [9, 17] of wireless networks to improve current wireless protocols and develop high performance, wireless-friendly applications.

Recent research has investigated WLAN performance when one host's traffic affects the performance of other hosts [2, 10]. This research has shown that when there is a WLAN host with weak wireless connectivity, the performance of all hosts can degrade considerably. These results are especially important for wireless home networks since, despite the relatively small size of most homes, the quality of wireless links in the home are not guaranteed, regardless of transmission power or rate [17]. However, the results thus far focus only on analyzing [10] and measuring performance at the wireless MAC layer [2], and thus they only provide meaningful results under narrow conditions. Previous work [13] indicates that performance aspects of the link layer, network layer and application layer can be inter-layer related. This suggests that effective models of infrastructure WLAN performance need to be aware of interactions between the network layers.

This paper provides insight into the performance interconnection of simultaneous applications running over the Internet to last-mile wireless infrastructure networks. Leveraging previously developed tools, experiments on a IEEE 802.11g WLAN network measure performance across the wireless link layer, network layer, transport layer, and application layer.

The contributions of this paper include: 1) confirmation of the performance anomaly modeled by Heusse *et al.* [10], whereby the 802.11 CSMA/CA channel access method causes a host with poor wireless connectivity to degrade the throughput of other hosts with better connectivity; 2) refinement of the results in [10], showing the anomaly dominates performance only when the wireless hosts use the same transport protocols, while heterogenous host protocols can push the performance bottleneck elsewhere; 3) confirmation of the results by Bai and Williamson [2], showing streaming over UDP to a host with poor wireless connectivity causes AP queue overflow that degrades the performance of other wireless hosts; 4) refinement of the results in [2], showing the AP queue overflow does not occur when streaming over TCP or when streaming UDP below the effective wireless capacity; and 5) demonstration that the behavior of the application

WMuNeP'05, October 13, 2005, Montreal, Quebec, Canada.

influences performance above and beyond performance predicted at the wireless and transport layer. The sum of these contributions illustrate the intertwined effects between network layers for 802.11 transmissions.

2. RELATED WORK

Understanding the performance of a flow traversing over a wired Internet environment to a wireless LAN has been the subject of many research papers. However, two aspects particularly germane to this study are the transport protocol's reaction to wireless losses and the interactions between two or more wireless hosts experiencing heterogeneous wireless transmission quality.

While much of the published research on wireless transport protocols involves TCP modifications (e.g. TCP-Westwood[1]) that alter TCP's reaction to wireless packet losses and MAC layer retries, this paper provides a multi-level view of the impact of TCP and UDP applications on hosts with poor wireless connectivity accessing the Internet through an access point (AP) used concurrently by other wireless hosts. Thus, this section considers only related work that focuses on capturing the interactions between wireless hosts.

Examples of earlier analytic models of IEEE 802.11 that capture detailed components of the wireless channel access mechanism include Cali *et al.* [5] and Bianchi [4]. Cali focuses on theoretical WLAN efficiency by dynamically determining the optimal local contention window size from the number of active wireless hosts and the average packet size. Bianchi extends this analytic model under ideal channel conditions to determine throughput limits for 802.11 with and without the RTS/CTS mechanism. Building on these two analytic models, more recent research efforts [6, 7, 15] include MAC layer retries and bit error rate in their models to determine delays and service times experienced by IP layer packets.

While these analytic models emphasize that wireless link layer contention impacts performance, they fail to account for significant aspects of newer wireless schemes such as 802.11b and 802.11g that can dynamically adapt the target wireless capacity to the host transmission quality. However, recent WLAN measurement studies provide new insights into these issues.

Heusse *et al.* [10] introduce the term *performance anomaly* to characterize the impact of a slow wireless host that transmits at a degraded target wireless capacity (e.g., 1 Mbps) compared to a fast wireless host that transmits at 11 Mbps (or 54 Mbps for 802.11g). Using a simplified version of a earlier analytic model to characterize the 802.11 backoff and MAC retry policies, the authors derive a channel contention-based result that claims the fast host's maximum throughput is degraded to the slow host's throughput. They conduct WLAN measurements that show moderate agreement with this result. While they assume degraded wireless capacities are actually due to bad transmission quality, their model and experiments both assume low bit error rates.

Bai and Williamson [2] measure the performance of two hosts streaming video over UDP through a common AP. Their results show that a mobile host streaming over UDP can suddenly enter a location with bad wireless connectivity, thereby seriously degrading the performance of another host streaming with good wireless conditions. They claim that the host in the bad location causes the UDP traffic to backlog since the wireless frames cannot be transmitted as fast as they arrive, causing the AP queue to overflow.

In a recent study, Yarvis *et al.* [17] examine characteristics of houses, physical location and wireless technology to show that home WLANs can be highly asymmetric and that transmission quality can vary significantly. Similar to the results of Aguayo *et al.* [1], they conclude that there is a low correlation between loss rate and distance and that precise node location is the single most important factor in determining the quality of wireless communication. Note, both of these studies ([17] and [1]) involve individual constant rate transmissions where all other wireless hosts are idle.

3. METHODOLOGY

This section discusses the experimental methodology used to investigate the multi-layered impact of a wireless client with bad connectivity on the performance of Internet traffic going to a wireless host with good connectivity through a common wireless access point.

3.1 Tools

For this investigation, several previously-developed [13] measurement tools for collecting data at multiple network layers were installed on two laptops. Table 1 lists the tools employed in this study and provides examples of the performance metrics available from each tool.

Table 1: Measurement Tools

Tools	Performance Metrics
UDP Ping	Round-Trip Time Packet Loss
Typeperf	Wireless Throughput Wireless Channel Capacity
WRAPI	Wireless Frame Retries Received Signal Strength Indicator (RSSI)

For network layer performance metrics such as round-trip time and packet loss rate along the flow path, an internally developed tool called *UDP ping* is used. Preliminary experiments revealed that since the standard ICMP *ping* provided by Windows XP waits for the previous ping reply or a timeout before sending out the next ping packet, a constant ping rate could not be maintained over poor wireless conditions where round-trip times longer than 3 seconds were recorded. Thus, a customized ping tool using UDP packets was built to provide constant ping rates, ping intervals configurable in milliseconds, and configurable ping packet sizes.

At the wireless data link layer, a publicly-available library, called *WRAPI* [3] was enhanced to collect information at the wireless streaming host that includes: signal strength, frame retransmission counts and failures, and information about the specific wireless access point (AP) that handles the wireless last hop to the host.

Additionally, *typeperf*, a performance monitoring tool built-in to Windows XP, is used to collect network data including received bitrate and the current wireless target capacity.

3.2 Experimental Design

Figure 1 shows the experimental setup. The wireless portion of the WPI campus network is partitioned from the wired infrastructure. Except for the last hop from the AP

[1]http://www.cs.ucla.edu/NRL/hpi/tcpw/

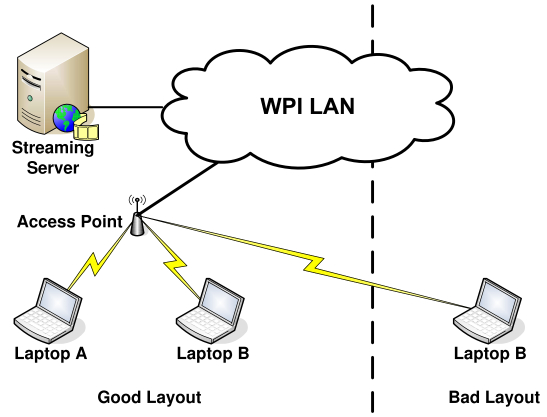

Figure 1: Measurement Setup

to the wireless hosts, all traffic traverses the same network path from a single server on the wired campus network to a common AP. The WPI wireless LAN uses Airespace[2] APs that provide IEEE 802.11g wireless service.

The two wireless hosts used in the experiments are labeled as host A and host B. Host A is a Toshiba laptop and host B is a Sony Vaio laptop. Both laptops run Windows XP home edition with Service Pack 1 and use Netgear WG511 802.11g network adaptors.

On each host, the UDP ping, typeperf and WRAPI tools described in Section 3.1 are run. Typeperf collects data every 1 second, WRAPI collects data every 500 ms, while UDP ping collects data on each 1350 byte packet sent every 200 ms.

Although these tools are deployed concurrently, baseline measurements indicated the tools together consume only about 3% of the processor time on either host and send only 5 packets per second. Given that the streaming videos consumed at least 35% of the processor time and AP beacons send 10 packets per second, the assumption is the measurement tools do not significantly impact wireless performance.

Since host mobility is not part of this study, all experiments involve one or two stationary laptops in one of two distinct configuration layouts at pre-determined locations in the WPI Computer Science Department building. In the first layout, both laptops are placed in locations that provide good wireless connectivity. In the second layout, host A remains at its good location while host B is placed at a location with bad wireless connectivity. Location identification and classification come from previous experiments [13] such that a good location has an average $SignalStrength \geq -70dBm$ and a bad location has an average $SignalStrength \leq -75dBm$.

To simplify coordination of concurrent flow measurement, a single server running Windows Server 2003 standard edition was used for all experiments. To verify the server processor was not a bottleneck, baseline experiments using two distinct servers, one for each laptop, were run. These results show that the two server and one server setups yield nearly identical throughput to the wireless LAN hosts (see [8] for details).

The offered load on the wireless LAN comes from two applications: the downloading of a large file and the streaming of a high-bandwidth multimedia clip. These two heavy-load applications were chosen to stress the wireless LAN such that congestion and channel contention would be observed and measured. *wget*, a publicly-available TCP download application,[3] is used to download a 400 MByte file from the server to a wireless host. Windows Streaming Media (v9.0), developed by Microsoft,[4] is used to stream a high-motion, 352×288 resolution, 24 frames per second, 2 minute video clip to a wireless host. The multimedia clip is encoded at a bitrate of 5.0 Mbps, with 4.8 Mbps for video and 0.2 Mbps for audio. The server is configured to support the two standard streaming transport protocol choices: TCP and UDP.

Table 2: Experiment Cases

Case	Host A	Host B	
	Good Location	Good Location	Bad Location
1	TCP Download	-	-
2	TCP Download	TCP Download	-
3	TCP Download	-	TCP Download
4	TCP Download	-	UDP Stream
5	TCP Download	-	TCP Stream
6	TCP Download	TCP Stream	-
7	TCP Download	UDP Stream	-
8	-	-	TCP Stream
9	-	-	UDP Stream

Table 2 lists the nine combinations (cases) of application workloads discussed in this paper. A dash in the table implies no application is running at that location. For example, case 4 represents a wireless measurement experiment where host A at a good signal location downloads a file from the server over TCP while host B at a bad signal location streams a video from the server over UDP.

At the beginning of each experimental run, the measurement tools described in Section 3.1 are started before the applications. To reduce the potential variability of the physical environment (as noted in [17]), the two laptops were placed in exactly the same locations with the same physical orientation for all the experiments. All experiments were conducted at night when no moving people were around and in locations known to have little wireless traffic in the evening. While each experiment produced about two minutes of usable performance data, only data between the 50th and 100th second are analyzed. This provides time to get beyond both the initial wireless experiment start up turbulence and the standard data rate burst used by streaming media players to quickly fill their playout buffers.

3.3 Consistency

Each of the nine cases in Table 2 were repeated three times to get some sense of the stability of the external environment and to guard against sporadic interference that might cause a particular run to yield inconsistent results. Figure 2 provides data from all three runs for case 4. The six graphs in the figure demonstrate that the performance patterns for the target wireless link capacity and the measured signal strength quality do not change significantly across the three runs. Comparing the results across multiple runs for the other eight cases yielded similar behavior. While there were a few cases where there was evidence of obvious interference, the length of the interference signal was short relative to the two minute video clip and/or file download. Thus, from the three runs the dominant performance characteristic was clearly discernable despite small-duration interference within a given run.

[2]http://www.airespace.com/

[3]http://www.gnu.org/software/wget/wget.html

[4]http://www.microsoft.com/windows/windowsmedia

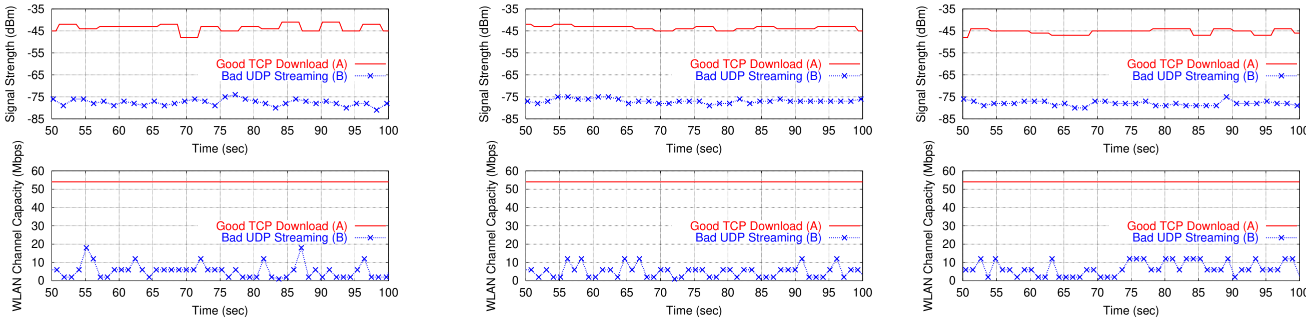

Figure 2: Wireless Signal Strength and Channel Capacity for Three Separate Runs

4. ANALYSIS

4.1 802.11 Performance Anomaly

The analysis begins with Figure 3 where the application layer throughputs for the first three cases in Table 2 are compared to show the impact of a host in a bad location downloading a file while a host in a good location concurrently downloads a file. For comparison, Figure 3(a) displays 50 seconds of measured throughput for case 1 where only host A (in a good location) is downloading a file while host B lies dormant. Note, the average throughput of 18.8 Mbps for the single host is significantly lower than the 54 Mbps maximum target capacity for an 802.11g channel but close to the maximum effective throughput calculated after overhead is taken into account, as in [12, 16].

Figure 3(b) shows throughput for case 2. With both downloads going to good locations, host A receives an average throughput of 9.3 Mbps and host B receives 9.6 Mbps. Thus, the two wireless hosts receive approximately half the throughput obtained by the single host at a good location. However, Figure 3(c) indicates that for case 3 the download to a bad location causes the throughput for both hosts to significantly degrade. In case 3, host A at a good location receives an average throughput of 2.8 Mbps while host B at a bad location receives an average throughput of only 2.1 Mbps.

Comparing the wireless signal strength (received signal strength indicator, or RSSI) in case 2 against the wireless signal strength case 3 in Figure 4 shows that the wireless signal for the host in the good location is not affected by the signal of the host in the bad location. This is reflected in the wireless target channel capacities in Figure 5, where the download to the good location consistently yields a target link capacity of 54 Mbps regardless of the location of host B. However, when host B is at a bad location its average target channel capacity falls below 11 Mbps.

The drop in the throughput of the host in the good location in case 3 is due to the IEEE 802.11 Distributed Coordination Function (DCF). Since DCF provides all hosts with an equal probability to access the wireless channel, hosts operating with a higher channel capacity wait nearly as long on average between sending packets as hosts operating at lower channel capacities. Thus, the average throughput of all hosts is reduced to the throughput of the host with the lowest channel capacity. These results are consistent with the anomaly discussed by Heusse *et al.* [10] and show their model of channel contention to be relevant even when 802.11g dynamically adapts the target channel capacity.

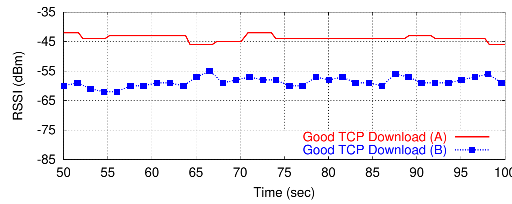

(a) Wireless RSSI for TCP Download in a Good Location and TCP Download in a Good Location

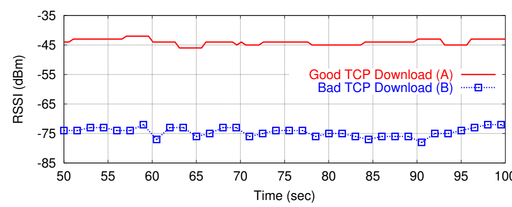

(b) Wireless RSSI for TCP Download in a Good Location and TCP Download in a Bad Location

Figure 4: Wireless Received Signal Strength Indicators

4.2 The Effect of the AP Queue

Next, the analysis investigates the difference in the 802.11 anomaly when the host at a bad location streams a multimedia clip using UDP instead of downloading a file using TCP.

Figure 6(b) graphs throughput for case 4 where host A at a good location downloads a file while host B at a bad location streams multimedia over UDP. Comparing this data with the results in Figure 6(a) shows that host A's throughput is essentially eliminated by the UDP stream coming to host B. Host A has a terrible average throughput of 0.3 Mbps while host B has an average throughput of 2.5 Mbps. The severe degradation in performance cannot be explained by the previously discussed 802.11 performance anomaly alone. The Heusse model [10] only accounts for throughput degradation caused by sharing the wireless capacity. The degraded TCP throughput for the host in the good location in case 4 may also be due to packet loss and higher round-trip times. Thus, loss and throughput at multiple network layers are now examined.

Figure 7 compares the wireless frame retries for case 3 (TCP download and TCP download) against wireless frame retries for case 4 (TCP download and UDP stream). Due

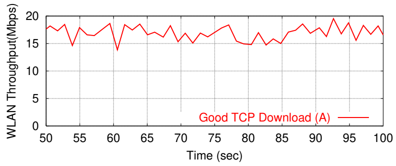

(a) TCP Download in a Good Location

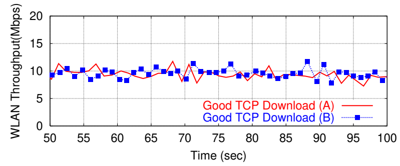

(b) TCP Download in a Good Location and TCP Download in a Good Location

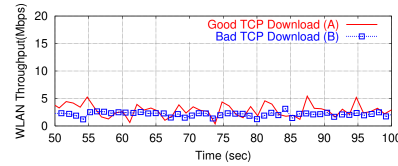

(c) TCP Download in a Good Location and TCP Download in a Bad Location

Figure 3: Throughput Comparison

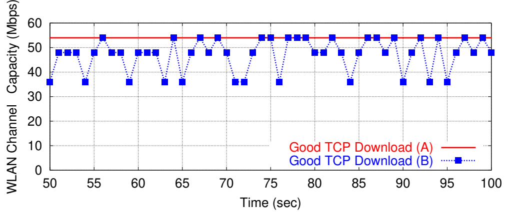

(a) Wireless Channel Capacity for TCP Download in a Good Location and TCP Download in a Good Location

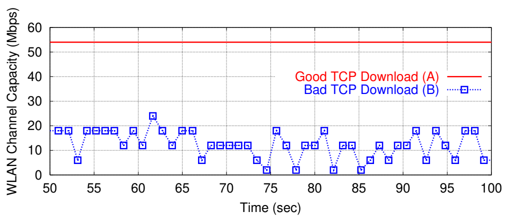

(b) Wireless Channel Capacity for TCP Download in a Good Location and TCP Download in a Bad Location

Figure 5: Wireless Channel Capacities

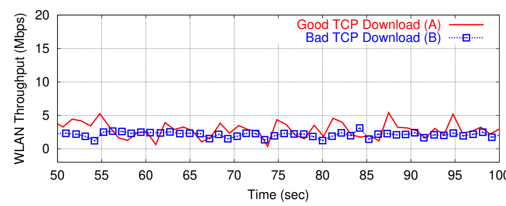

(a) TCP Download in a Good Location and TCP Download in a Bad Location

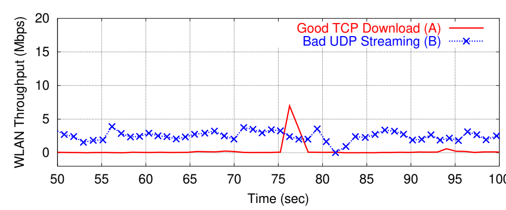

(b) TCP Download in a Good Location and UDP Streaming in a Bad Location

Figure 6: Throughput Comparison

to the bad wireless conditions, host B in Figure 7(a) records an average retry fraction of about 0.2, while host A at the good location has a retry fraction of only about 0.05. In Figure 7(b), both hosts experience bursty frame retry behavior. Wireless retry behavior alone cannot explain the performance difference between case 3 and case 4. These results suggest refinement of the results in [10], showing the anomaly dominates performance only when the wireless hosts use the same transport protocols, while heterogenous host protocols can push the performance bottleneck elsewhere. To get more insight into the recorded performance, it is necessary to also consider packet behavior at higher layers.

IEEE 802.11 MAC protocols respond to bit errors or frame contention losses by retransmitting frames up to a specified retry limit. The wireless MAC layer thus insulates the IP layer above from packet losses caused by bad wireless conditions, except when the retry limit is exceeded and the dependent IP layer packet is dropped. Comparing the retry fraction in Figure 7(a) to the IP (UDP ping) packet loss rate in Figure 8(a) demonstrates this effect. Host B in a bad location has many wireless frames retransmitted, but the UDP packet loss is near zero, comparable to that of host A.

However, Figure 8(b) presents a completely different picture when host B is in a bad location and streams multimedia over UDP. The extremely high UDP ping packet loss rates shown in Figure 8(b) are not explained by the 802.11 anomaly, but rather are due to network layer congestion at the AP queue. The 5 Mbps UDP stream is unresponsive to the limited wireless capacity and overflows the AP queue, continuing unabated in the face of extreme packet loss. Thus, the difference in UDP ping loss behavior between Figure 8(a) and Figure 8(b) shows how the 802.11 anomaly model does not capture congestion in the AP queue.

Figure 9 presents cumulative distribution functions (CDFs) for UDP ping round-trip times concurrently sent from both wireless hosts for cases 3 and 4. Figure 9(a) clearly shows higher round-trip times for the host downloading at a bad location compared to the host downloading at a good location. This moderate increase in round-trip time can be attributed to the increase in wireless layer retry fraction seen in Figure 7(a).

Figure 9(b) shows cumulative distribution functions of the round-trip times for case 4. Note the x-axis range in Figure 9(b) is considerably larger than the x-axis range in Figure 9(a). While the left side of the round-trip time distributions in case 4 are different from those in case 3, it is difficult to draw conclusions from these cumulative distributions

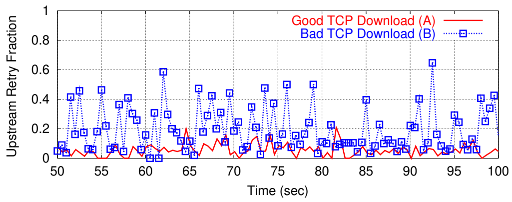

(a) Wireless Layer Retry Fraction for TCP Download in a Good Location and TCP Download in a Bad Location

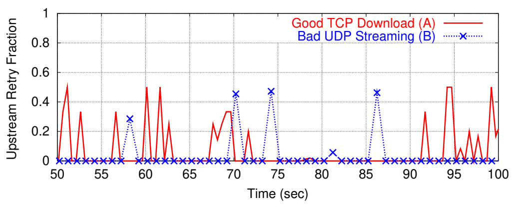

(b) Wireless Layer Retry Fraction for TCP Download in a Good Location and UDP Streaming in a Bad Location

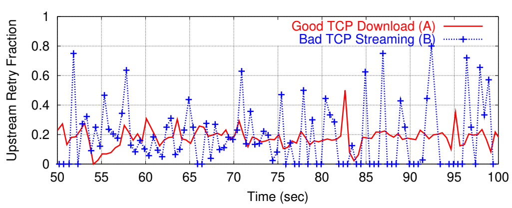

(c) Wireless Layer Retry Fraction of for TCP Download in a Good Location and TCP Streaming in a Bad Location

Figure 7: Wireless Layer Retry Fraction Comparison

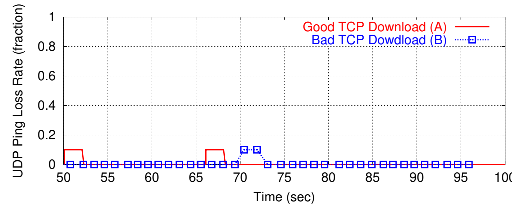

(a) UDP Ping Packet Loss Fraction for TCP Download in a Good Location and TCP Download in a Bad Location

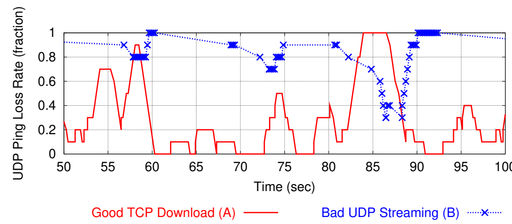

(b) UDP Ping Packet Loss Fraction for TCP Download in a Good Location and UDP Streaming in a Bad Location

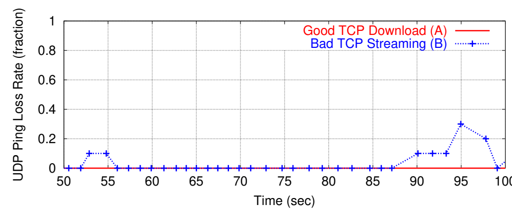

(c) UDP Ping Packet Loss Fraction for TCP Download in a Good Location and TCP Streaming in a Bad Location

Figure 8: Packet Loss Fraction Comparison

alone because as Figure 8(b) has already shown, the UDP packet loss rates are very high. Thus, the data points are sparse for the tail of the CDFs when there is a host streaming UDP from a bad location. However, the high round-trip times shown in Figure 9(b) provide further evidence that both flows encounter a large AP queue.

These results are consistent with the results from [2] and show streaming multimedia over UDP to a host with poor wireless connectivity causes the AP queue to overflow, degrading the performance of all wireless hosts. However, when UDP streaming is replaced with TCP streaming, the AP queue is not the bottleneck. Comparing Figures 7(c), 8(c) and 9(c) with their corresponding UDP streaming figures shows the AP queue does not fill up at all with TCP streaming. This suggests a refinement of the results in [2], showing the AP queue overflow does not occur when streaming over TCP or when streaming UDP below the effective wireless capacity. The impact on the host in the good location is caused by the intertwining of effects of the lower wireless layer, as shown in Section 4.1, and the upper application layer, as shown in the next Section.

4.3 The Effect of the Application Layer

Figure 10 provides a top-level view of the intertwined effect of the application and the node location on application layer throughput. In Figures 10(a) and 10(b) where both hosts are at good locations, host B is able to stream the 5 Mbps encoded bitrate over both TCP and UDP. Note, the UDP stream actually leaves slightly more capacity for the concurrent TCP download than does the TCP stream.

Figures 6 and 10(c) can be compared to see the effect of TCP versus UDP intertwined with whether or not the application is a file download or streaming multimedia while the host is at a bad location. The difference in throughput for the TCP download versus the TCP stream at a bad location is because streaming media servers can react to indications of inadequate available capacity by performing media scaling at the application layer, thereby reducing the streaming bitrate. This effect can be seen in Figure 10(c) where the media server scales down the streaming bitrate around time 70. The TCP stream in a bad location yields a throughput lower than the TCP download in bad location in Figure 6(a). However, both of these TCP-based applications leave some available capacity for the concurrent TCP download in a good location. This lies in stark contrast to the previously analyzed UDP stream in a bad location in Figure 6(b) that wipes out throughput for the TCP download in the good location.

One more set of measurements is provided in Figure 11

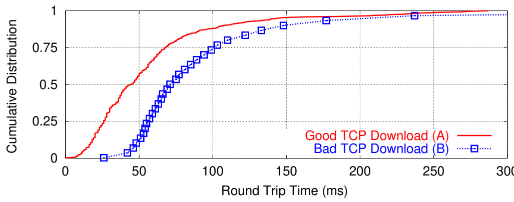

(a) CDF of Round-Trip Time of TCP Download in a Good Location and TCP Download in a Bad Location

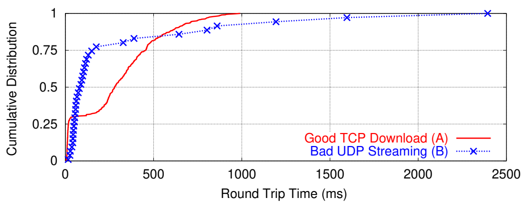

(b) CDF of Round-Trip Time of TCP Download in a Good Location and UDP Streaming in a Bad Location

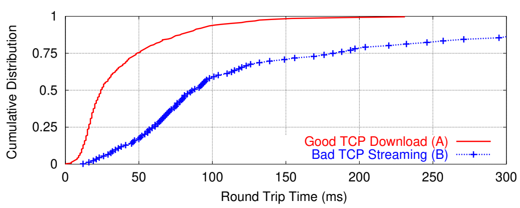

(c) CDF of Round-Trip Time of TCP Download in a Good Location and TCP Streaming in a Bad Location

Figure 9: Round-Trip Time Comparison

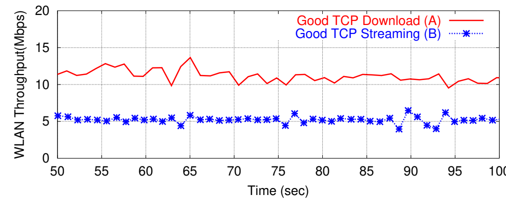

(a) Throughput of TCP Download in a Good Location and TCP Streaming in a Good Location

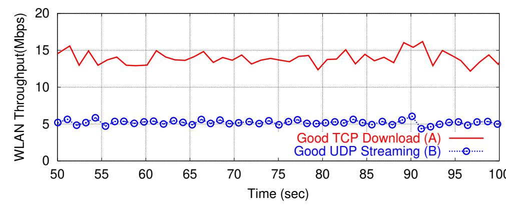

(b) Throughput of TCP Download in a Good Location and UDP Streaming in a Good Location

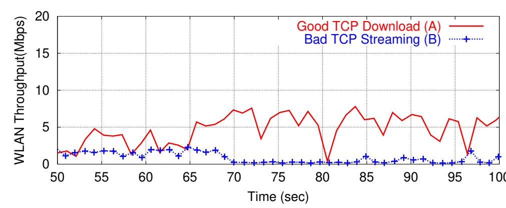

(c) Throughput of TCP Download in a Good Location and TCP Streaming in a Bad Location

Figure 10: Throughput Comparison

for cases 8 and 9. This is another attempt to separate out the intertwined effects of the streaming application from the choice of TCP or UDP for the transport protocol. For these cases, the hosts streaming over TCP (case 8) or UDP (case 9) are in bad locations, but do not contend with heavy-bandwidth applications in good locations. With almost no channel or AP queue contention due to a concurrent flow, the general shapes of the application throughput for TCP and UDP are quite similar in Figures 11(a) and 11(b). The initial throughput spikes lasting until about 20 seconds in both graphs are indicative of the streaming servers initially attempting to send at a high data rate to fill the media player playout buffer, consistent with results from [14]. The later ebbs and flows in both throughput time lines can be explained by attempts by the Windows Media server to match the encoded streaming bitrate to the available capacity. When the media players playout buffer drains, the server reduces the encoded bitrate in an attempt to "thin" the stream to match the observed capacity. At other times, the server deems there is additional capacity available and increases the encoded bitrate in an attempt to provide better quality. The critical point is that the performance is not simply explained by the host being at a bad location or by the transport protocol chosen, alone.

5. CONCLUSION

This paper takes a step towards providing a better understanding of 802.11 networks under a typical home network condition for accessing the Internet, namely when one host has good wireless connectivity while another host has bad wireless connectivity. Carefully designed experiments with previously developed tools induce heavy load on a real 802.11g wireless network, allowing measurements at multiple network layers. The tools capture the performance at the wireless link, network, transport and application layers and enable analysis of the intertwined effects between network layers for 802.11g transmissions.

Our experiments demonstrate that multiple 802.11g conversations sent through a common wireless access point (AP) cause channel contention that lowers effective throughput. By varying the traffic from the higher layers in this investigation, our experiments provide the following observations beyond the result that wireless link layer contention impacts performance:

1. *Network layer queues at the wireless access point impact performance.* When the wired network layer throughput is higher than the effective capacity at the wireless link, the access point queues can severely lower performance for all flows traversing the AP due to increased queuing delay and buffer overflow. Performance degra-

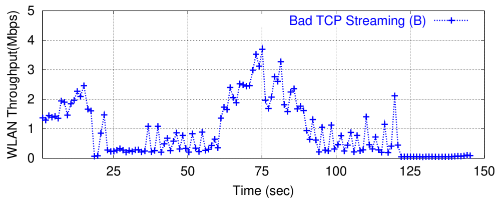

(a) TCP Streaming in a Bad Location

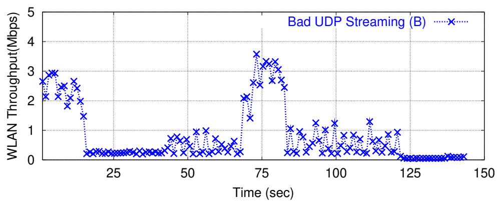

(b) UDP Streaming in a Bad Location

Figure 11: Light Load Throughput Comparison

dation for all wireless clients is exacerbated when one client has bad wireless connectivity.

2. *The choice of transport protocol impacts performance.* TCP and UDP clients at bad locations affect good clients differently due not only to wireless channel contention but also due to contention for the AP buffer. TCP flows self-contend with their own acknowledgments and unresponsive UDP flows are more likely to overflow the AP queue.

3. *Application layer behavior also impacts wireless performance.* Above the transport layer, Internet applications may adjust to the wireless network environment. For example, while a file download relies on TCP to adapt to low quality wireless conditions, a streaming media server will invoke media scaling in an attempt to stream at an encoded bitrate that is below the perceived available streaming capacity.

The significance of the above observations can be seen in recent work by Yoo *et al.* [11] that proposes to adjust the wireless frame size proportionally to the available wireless capacity. While this methodology does address the 802.11 performance anomaly, it does not address the other intertwined effects caused by higher layer protocols and applications running at low quality wireless locations.

Future work includes building an analytic model that encompasses wireless contention, the AP queue, and an application, representing the intertwined effects shown in this paper's measurements. As a starting point, the application can likely be modeled as a bulk download, but there in also an opportunity for more sophisticated models of streaming media or Web browsing applications.

While the performance measurements in this paper are from a real 802.11 network, the effects of cross traffic such as AP beacons and WLAN traffic to other APs were not controlled. Understanding the effects of interfering, competing or contending traffic will provide additional insights into wireless network performance.

6. REFERENCES

[1] D. Aguayo, J. Bicket, S. Biswas, G. Judd, and R. Morris. Link-level Measurements from an 802.11b Mesh Network. In *Proceedings of ACM SIGCOMM*, Portland, OR, USA, Sept. 2004.

[2] G. Bai and C. Williamson. The Effects of Mobility on Wireless Media Streaming Performance. In *Proceedings of Wireless Networks and Emerging Technologies (WNET)*, pages 596–601, July 2004.

[3] A. Balachandran and G. Voelker. WRAPI – Real-time Monitoring and Control of an 802.11 Wireless LAN. Technical report, CS at UCSD, 2004.

[4] G. Bianchi. Performance Analysis of the IEEE 802.11 Distributed Coordination Function. *IEEE Journal on Selected Areas in Communications, Wireless Series*, 18(3):535–547, Mar. 2000.

[5] F. Calì, M. Conti, and E. Gregori. IEEE 802.11 Wireless LAN: Capacity Analysis and Protocol Enahncement. In *Proceedings of IEEE INFOCOM*, pages 142–149, San Francisco, CA, USA, Mar. 1998.

[6] M. Carvalho and J. Garcia-Luna-Aceves. Delay Analysis of IEEE 802.11 in Single-Hop Networks. In *Proceedings of IEEE International Conference on Network Protocols (ICNP)*, Atlanta, Georgia, USA, Nov. 2003.

[7] P. Chatzimisios, V. Vitsas, and A. C. Boucouvalas. IEEE 802.11: Packet Delay - A Finite Retry Limit Analysis. In *Proceedings of IEEE Globecom 2003*, pages 950–954, San Francisco, CA, USA, 2003.

[8] J. Gretarsson, F. Li, M. Li, A. Samant, H. Wu, M. Claypool, and R. Kinicki. Performance Analysis of the Intertwined Effects between Network Layers for 802.11g Transmissions. Technical Report WPI-CS-TR-05-09, CS Department, Worcester Polytechnic Institute, May 2005.

[9] T. Henderson, D. Kotz, and I. Abyzov. The Changing Usage of a Mature Campus-wide Wireless Network. In *Proceedings of International Conference on Mobile Computing and Networking (MobiCom)*, Philadelphia, PA, USA, Sept. 2004.

[10] M. Heusse, F. Rousseau, G. Berger-Sabbatel, and A. Duda. Performance Anomaly of 802.11b. In *Proceedings of IEEE INFOCOM*, 2003.

[11] S. hwan Yoo, J.-H. Choi, J.-H. Hwang, and C. Yoo. Eliminating the Performance Anomaly of 802.11b. In *International Conference on Networking*, Apr. 2005.

[12] J. Jun, P. Peddabachagari, and M. Sichitiu. Theoretical maximum throughput of IEEE 802.11 and its applications. In *Proceedings of the 2nd IEEE International Symposium on Network Computing and Applications (NCA'03)*, pages 249–256, Cambridge, MA, USA, Apr. 2003.

[13] F. Li, J. Chung, M. Li, H. Wu, M. Claypool, and R. Kinicki. Application, Network and Link Layer Measurements of Streaming Video over a Wireless Campus Network. In *Proceedings of the 6th Passive and Active Measurement Workshop (PAM)*, Boston, Massachusetts, USA, Apr. 2005.

[14] J. Nichols, M. Claypool, R. Kinicki, and M. Li. Measurements of the Congestion Responsiveness of Windows Streaming Media. In *Proceedings of ACM NOSSDAV*, June 2004.

[15] H. Wu, Y. Peng, K. Long, S. Cheng, and J. Ma. Performance of Reliable Transport Protocol over IEEE 802.11 Wireless LAN: Analysis and Enhancement. In *Proceedings of IEEE INFOCOM*, pages 599–607, New York, NY, USA, June 2002.

[16] Y. Xiao and J. Rosdahl. Throughput and delay limits of IEEE 802.11. *IEEE Communications Letters*, 6(8):355–357, Aug. 2002.

[17] M. Yarvis, K. Papagiannaki, and W. S. Conner. Characterization of 802.11 Wireless Networks in the Home. In *Proceedings of 1st workshop on Wireless Network Measurements (WiNMee)*, Riva del Garda, Italy, Apr. 2005.