THE IMPORTANCE OF AIR TIME ALLOCATION IN WI-FI QUALITY OF SERVICE

Eli Baruch
ARRIS

Abstract

As Cable and Telco service providers strive to own the subscriber's home experience by deploying ever more sophisticated Wireless Gateways and multi-service offerings; wireless technology plays a key enabling role in the way consumers use the various services. With the proliferation of Multi-SSID Wireless gateways that offer multiple services such as an in-home private network, Home Security and Appliance Monitoring, Public Hotspots and Video over Wi-Fi; the need to monitor and assure basic service levels and an overall Quality of Service (QoS) becomes essential to the user's quality of experience.

Through a series of real-life tests, this paper will show how these competing services can impact the home user's wireless experience, in both single user and multi user environment. Also demonstrated, is how the shortcomings of existing wireless systems, the limits of existing WMM, and the lack of a proper QoS monitoring and enforcement mechanisms can degrade the user experience and create severe subscriber retention problems. This paper will show that the critical resource on the wireless network is not raw bandwidth management but rather comprehensive management of the time allotted on the air interface, or Air Time Management. Lastly a variety of options will be discussed illustrating methods of avoiding delivery of a poor user experience and guaranteeing an acceptable basic service level.

WIRELESS QUALITY OF SERVICE

Wireless 802.11 technologies were designed as a method of extending LAN-type service over the air. As such, they were seen as an extension of Ethernet LAN services, for which QoS usually did not play a major role in the end user's experience. Due to the nature of the wireless medium, and in anticipation of multiple types of traffic using the air interface, a basic QoS mechanism was defined in the 802.11e standard and adopted by the Wi-Fi Alliance as part of their certification and interoperability program under the name of Wi-Fi Multimedia (WMM)(1).

WMM defines four Access Categories (AC):

- VoIP—Very low throughput with highest priority and strict latency requirements
- Video—High priority with latency requirements
- Best Effort—Low priority
- Background—Lowest priority

Incoming traffic is tagged and assigned one of four different priorities. Individual packets are then directed to one of four internal queues and prioritized according to the AC into which they fall. Packets from higher priority ACs are transmitted with a smaller inter-frame space and a smaller random back-off window, which allows transmission to the wireless medium with less delay on average.

The WMM mechanism thus provides statistical priority for winning access to the air interface. The algorithm used by a WMM enabled Access Point is probabilistic and depends on two timing parameters that vary for each AC:

1. The minimum interframe space or Arbitrary InterFrame Space Number (AIFSN), and
2. The Contention Window (CW), a random backoff wait time.

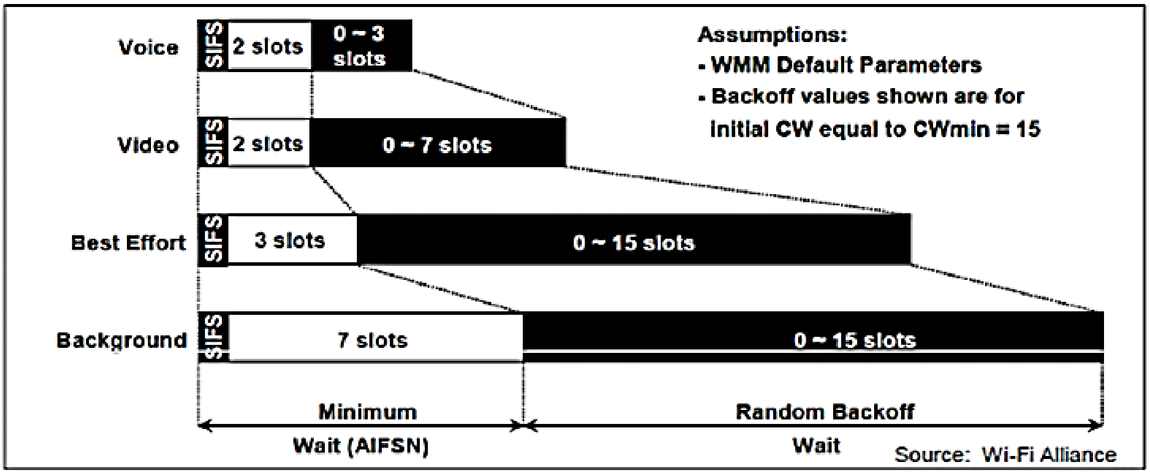

Figure 1: Wireless Multimedia – wireless QoS mechanism

After each collision, the CW is doubled until a maximum value, which also depends on the AC, is reached. As frames with the highest AC tend to have the lowest backoff value, they are more likely to be transmitted. However, little consideration was given to a multi-SSID scenario, in which different services may use different SSIDs to indicate the expectation of different levels of service. For example, a home security service may use a very limited amount of bandwidth. The user, however, expects that bandwidth to be available whenever needed. On the other hand, a hotspot service may have inconsistent bandwidth requirements, but the user does not expect such service to degrade or compete with the wireless home network's bandwidth.

Level of Service Expectations

A basic minimum level of service is expected from any Internet service provider. To fulfill this expectation, most service providers manage their High Speed Internet services with a dynamic set of DOCSIS service flows or using IP Differentiated Services protocols. These mechanisms govern how bandwidth is allocated to different services and users. Such schemes distribute the maximum overall bandwidth to and from the home by dynamically allocating bandwidth between preferred and best-effort services. The DOCSIS service flow mechanism, however, only governs the traffic allocation to and from the cable modem or wireless gateway. This mechanism does not enforce any priority over the wireless network. A fundamental question to be addressed is the actual workings of the intuitive picture depicted in Figure 2: Theoretical behavior of Dynamic Service Flow.

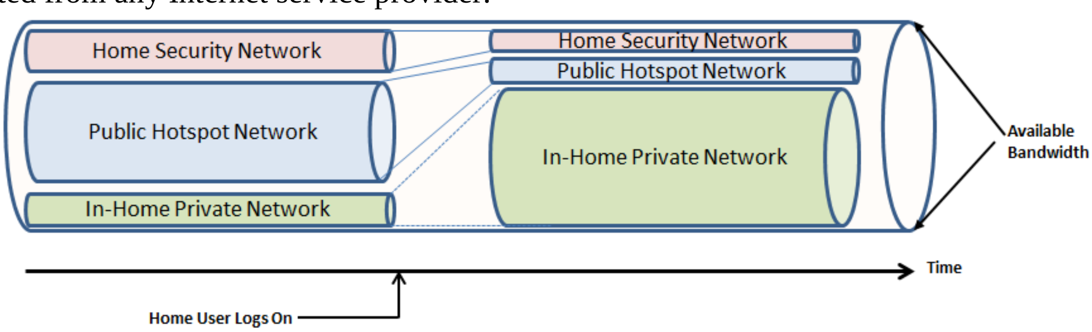

Figure 2: Theoretical behavior of Dynamic Service Flow

UNDERSTANDING AIR TIME

As a wireless signal propagates through the air, its signal strength and the ratio of signal to noise as seen by the receiving end diminishes as a function of distance, attenuation (e.g. obstacles , Multipath, and reflections), and temporal or other interference. The minimum received signal power level required to achieve a sufficient signal-to-noise ratio (SNR) is called "receive sensitivity." If the received signal power level falls below the receive sensitivity required for a given data rate, communication at that data rate becomes unreliable. A common manifestation of such unreliable communication is video stream buffering or freezing, as well as dual-band switchable clients (e.g. iPad) switching between 5GHz and 2.4GHz bands.

In order to maintain a reasonable SNR, the transmitter will modify the Modulation and Coding Scheme (MCS) by changing the modulation profile, the error correction scheme, and the number of spatial streams that send traffic to the receiver. Without going into a deep technical description, we can simply state the higher the SNR, the better the performance. Higher order QAM are not as robust against noise or other degradations as lower-order QAM. Using higher-order QAM without increasing the bit error rate requires a higher signal-to-noise ratio (SNR). The table in
Figure 3: 802.11n MCS index and data rate
lists the different Modulation and Coding Scheme index used by a 3x3:3 802.11n wireless device.

MCS	Modulation	Coding	Spatial Streams	Theoretical Data Rate (Mbps) GI = 800ns		MCS	Modulation	Coding	Spatial Streams	Theoretical Data Rate (Mbps) GI = 800ns	
				20MHz	40MHz					20MHz	40MHz
0	BPSK	1/2	1	6.5	13.5	12	16-QAM	3/4	2	78.0	162.0
1	QPSK	1/2	1	13.0	27.0	13	64-QAM	2/3	2	104.0	216.0
2	QPSK	3/4	1	19.5	40.5	14	64-QAM	3/4	2	117.0	243.0
3	16-QAM	1/2	1	26.0	54.0	15	64-QAM	5/6	2	130.0	270.0
4	16-QAM	3/4	1	39.0	81.0	16	BPSK	1/2	3	19.5	40.5
5	64-QAM	2/3	1	52.0	108.0	17	QPSK	1/2	3	39.0	81.0
6	64-QAM	3/4	1	58.5	121.5	18	QPSK	3/4	3	58.5	121.5
7	64-QAM	5/6	1	65.0	135.0	19	16-QAM	1/2	3	78.0	162.0
8	BPSK	1/2	2	13.0	27.0	20	16-QAM	3/4	3	117.0	243.0
9	QPSK	1/2	2	26.0	54.0	21	64-QAM	2/3	3	156.0	324.0
10	QPSK	3/4	2	39.0	81.0	22	64-QAM	3/4	3	175.5	364.5
11	16-QAM	1/2	2	52.0	108.0	23	64-QAM	5/6	3	195.0	405.0

Figure 3: 802.11n MCS index and data rate

A wireless gateway/access point will change the MCS used to communicate with different clients. In other words, the same client, depending on its location relative to the access point, may use a different MCS setting and therefore have a different level of data throughput at different locations. The further the client is from the access point; the fewer number of bits it will transmit in a given amount of Air Time. Moreover, the more time is required for communicating with one client, means that less time—and hence less overall bandwidth—remains for others using the same access point. It is important to note that the MCS index by which the access point will communicate with a given client is not fixed. As a wireless environment changes, an access point may switch MCS while transmitting to the same client based on SNR and Packet Error Rate (PER).

In order to illustrate the receive sensitivity distribution based on location, some use a color coded map, also known as a Heat Map, that shows the receive levels seen by a client from the access point. Such maps, although not always a direct representation of the bandwidth one might expect at different locations, give some indication of the behavior to expect from close and remote clients.

Figure 4: Example of receive levels “heat map”

The Impact of Location on Performance

To demonstrate the effect that a lower MCS due to distance has on the amount of Air Time that a client uses, we conducted several tests in a test house. The test house layout is depicted in**Error! Reference source not found.**.

We measured the traffic sent from the Access Point (AP) located at the main floor to a single client in location 9, the reception room on the main floor. We placed the client in location 9 10 feet away from the AP and in line of sight.

We also measured the traffic sent from the AP to location 6, the bathroom on the upper floor located at the edge of the house. Location 6 was 55 feet and 5 walls away from the AP.

We measured the throughput of downstream traffic from the AP to the client and the response time that the TCP connection needed to acknowledge the transfer of a fixed block of TCP packets. The AP we used was an ARRIS Wireless Gateway with Dual Band Concurrent 3x3:3 802.11N radios. We conducted all measurements using the 2.4GHz band.

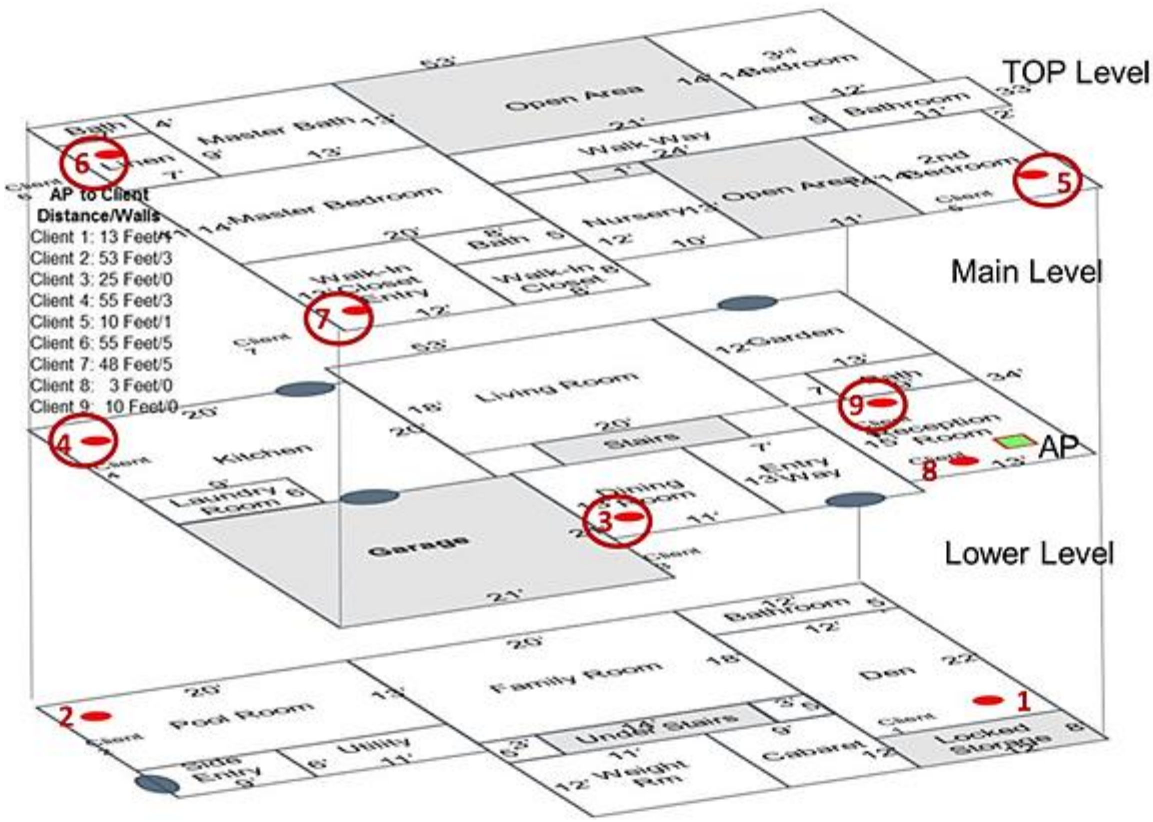

Figure 5: Wi-Fi test house client location map

The client in location 9 is close to the AP, which used MCS=22 to communicate with it. As result, it enjoyed good connection speeds, averaging 96.37 Mb/s, and fast response times, averaging of 0.834 sec.

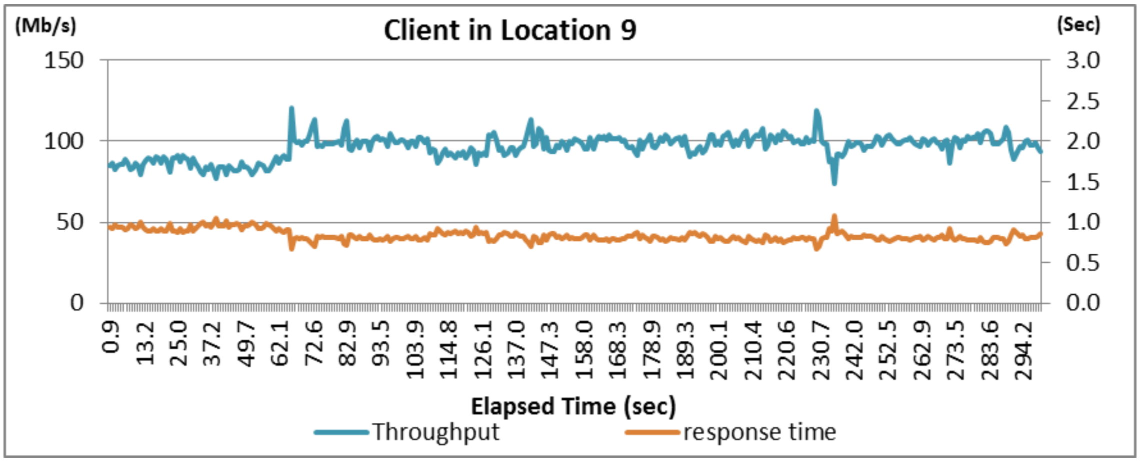

Figure 6: Throughput and response time measured in location

The client in location 6 started temporarily with MCS=17, continued for a while with MCS=19, and finished with MCS=20. As result, it experienced lower throughput than Client 9, with an average speed of 58.75Mb/s, and experienced longer response times, averaging 1.378 sec. In other words, in order to send the same amount of traffic, the client in location 6 used **40%** more time to complete the data transfer.

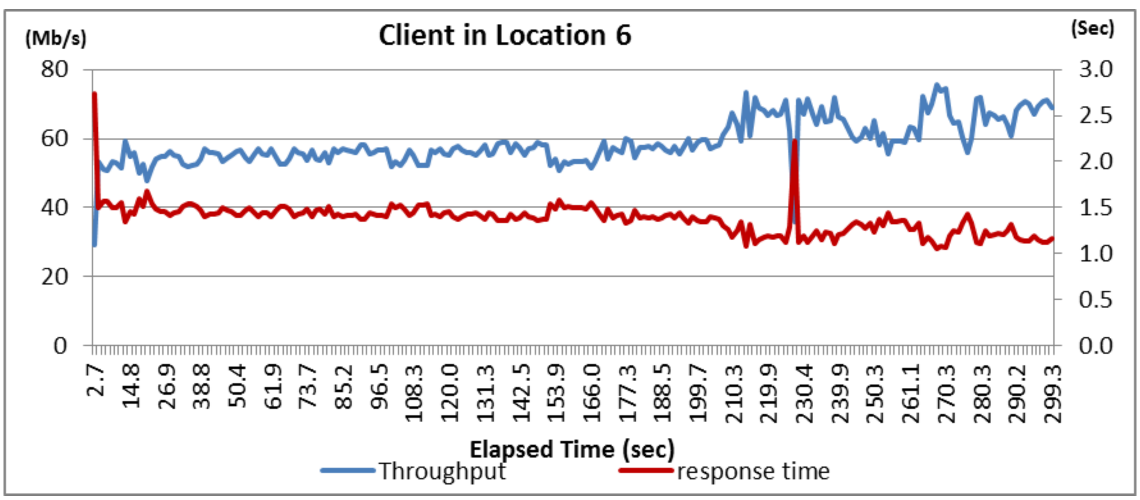

Figure 7: Throughput and response time measured in location 6

MULTI-CLIENT ENVIRONMENT

Our statistics show that, on average, there are six active wireless devices in a home that share access to the home wireless network **Error! Reference source not found.**. Since those different clients may reside in different rooms in the house, and at different distances from the AP, the overall aggregated available bandwidth in a multi-client environment is reduced slightly in comparison with a single device.

Assuming a typical package of 50Mb/s WAN-side Internet service, here is an example of test results for bandwidth usage and bandwidth distribution from amongst one to six wirelessly connected 802.11n devices in the home. In the example shown in Figure 8: Bandwidth usage in multi-client environment , all of the devices are located at similar distances from the house's main AP.
Note that even at similar distances and using the 802.11n standard, without any slower 802.11g or 802.11b clients on the network, there may still be differences in overall bandwidth between devices because of differences in their wireless characteristics (e.g., the number of receive and transmit antennas, power levels, etc.). The overall aggregated bandwidth used by the varying number of devices, however, is affected only slightly by the number of active devices. As one would expect, the more devices that are actively connected, the lower the overall bandwidth shared among all devices—including the individual bandwidth allocated per device.

Although when multiple clients are attached to the same AP, the overall available time on the wireless medium for data transmission is also reduced compared to a single device. The way the access point choses to distribute the Air Time significantly impacts the individual client data throughput and may as well impact the overall aggregated data throughput. In some cases, an AP may use a "fairness" algorithm to allow for the aggregated bandwidth to be allocated fairly between the clients. Such "fairness" algorithms may use a packet base round robin between clients. The drawback of this approach is that with TCP type of IP connection the client using the higher rate MCS will end up consuming the majority of the available bandwidth to the point that it may starve the other clients. Other algorithms may try to take into account time-based information such as Wi-Fi frames retransmissions and MCS rate.

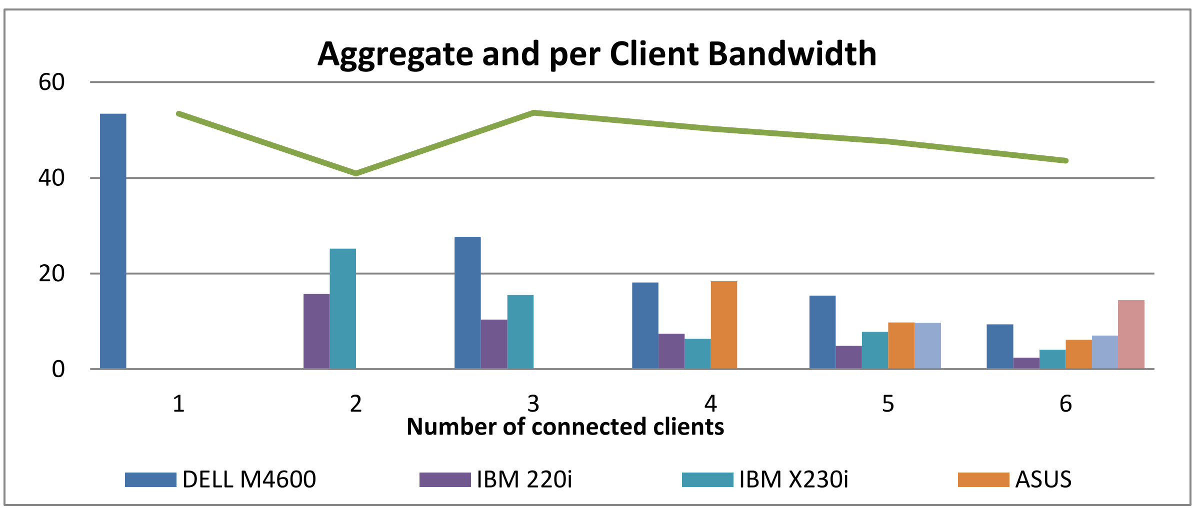

Figure 8: Bandwidth usage in multi-client environment

The example below shows test results of 10 clients connected to two different access points. Each access point uses a different algorithm to allocate Air Time between clients. Although the aggregated bandwidth consumed by the 10 clients in both tests was almost identical, 36Mb/s, note how differences in Air Time "fairness" algorithm manifests in each client's bandwidth.

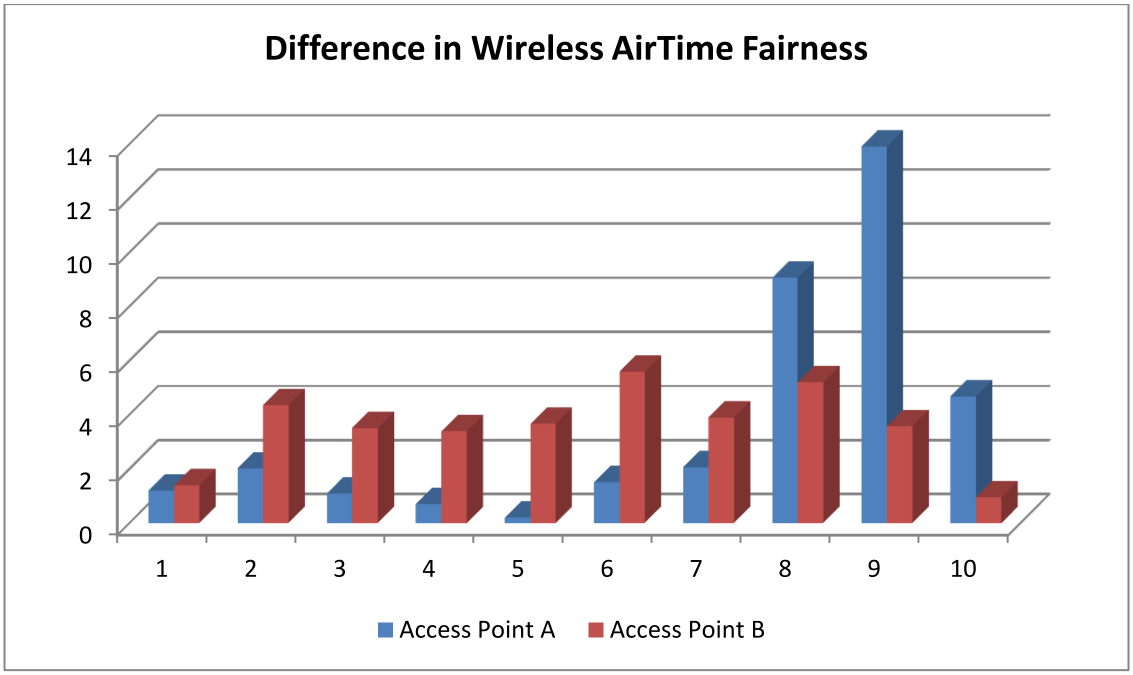

Figure 9: Bandwidth usage in multi-client environment

In the example shown in Figure 110 and Figure 121, six clients were placed at locations 3, 4, 5, 6, 7, and 9 in the test house. In order to eliminate the variability of the clients' Wi-Fi characteristics, the same version of client was used in each location. The AP was an ARRIS Wireless Gateway with a 2x2:2 2.4GHz 802.11N radio.

Although the six clients were located in different rooms, the results show a relatively even distribution of overall data throughput between the different clients, with the three clients closest to the AP—3, 5 and 9—enjoying higher throughput. As one can see from the graphs below, the bandwidth per client varies over time.

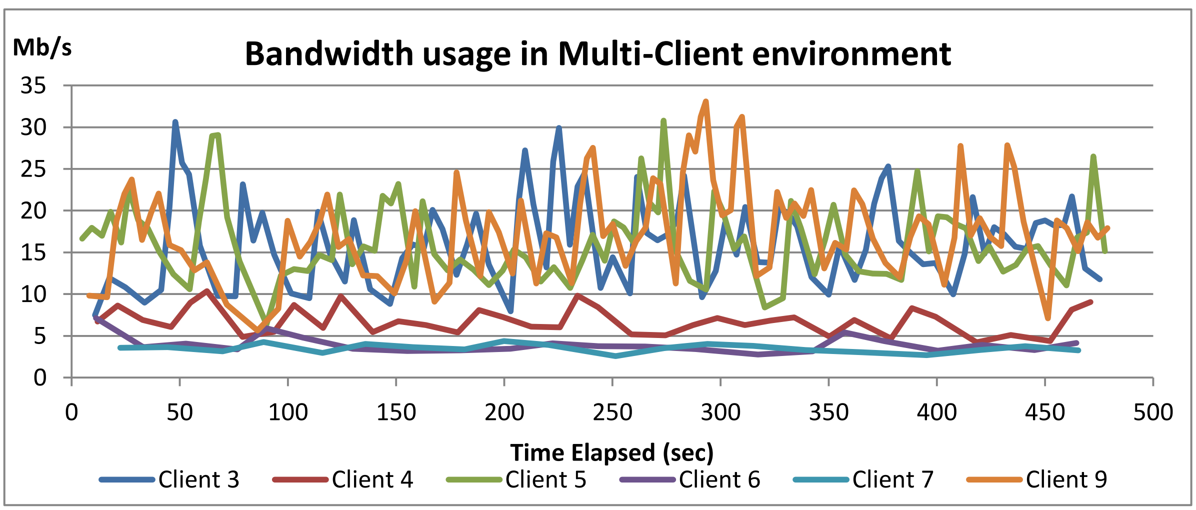

Figure 10: Bandwidth usage in multi-client environment

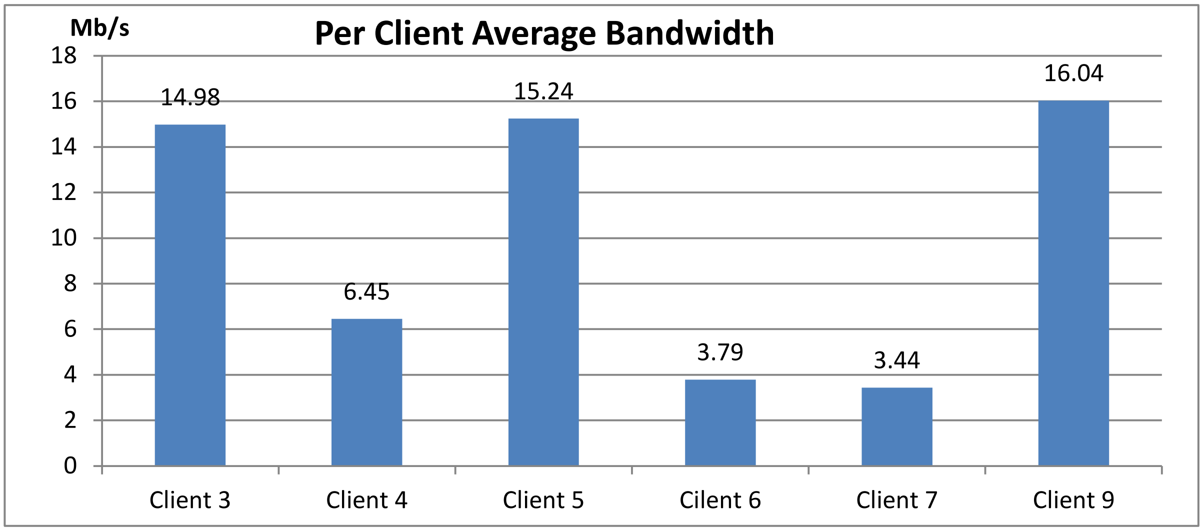

Figure 11: Average bandwidth usage in multi-client environment

	Location 3	Location 4	Location 5	Location 6	Location 7	Location 9
Average Bandwidth (Mb/s)	14.98	6.45	15.24	3.79	3.44	16.04
Average Response Time (Sec)	5.34	12.40	5.25	21.11	23.27	4.99
Percentage of Air Time Used	7.38%	17.14%	7.26%	29.16%	32.16%	6.90%

Figure 12: Air Time distribution in multi-client

If we take a closer look at the over-the-air allocation of time, the remote clients with lower MCS values take significantly more Air Time to deliver the same number of byte

THE IMPACT OF WIRELESS HOTSPOT SERVICE

Most service providers want to augment new and existing outdoor wireless deployments with home hotspot services using always-on, multi-SSID wireless gateways. The active hotspot serves Wi-Fi roaming users and/or offloads cellular traffic from the 3G/4G network onto the Wireless network and into the service provider's DOCSIS or IP backhaul networks.

Although some of the users connected to the hotspot SSID may be houseguests, they are more commonly roamers who are situated outside the house, but still within the wireless coverage range of the home access point. By the very nature of being outside, those roaming devices may be at the fringes of the home's wireless coverage and therefore would use a lower speed MCS compared to devices in the home.

As the number of roaming clients connected to the home hotspot increases, the impact on the existing clients in the home becomes more noticeable. These distant clients use a large portion of Air Time, which causes the in-home clients' bandwidth to decrease noticeably.

In order to demonstrate the impact of home hotspot roamers on in-home clients, the bandwidth and response time of Client 9 were monitored after situating two Wi-Fi clients just outside the house. These clients connected to the home's hotspot service and streamed an HD YouTube video. The home hotspot clients started streaming the movies approximately four minutes after test began.

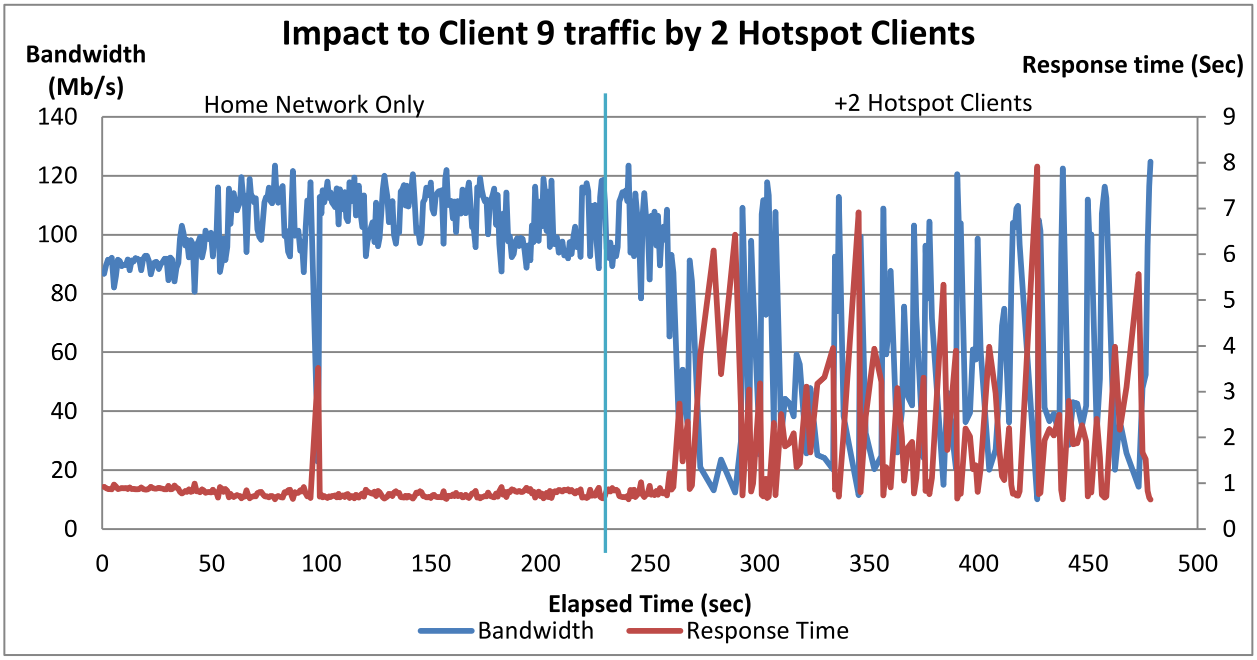

Figure 13: Client 9 bandwidth and response time
(Hotspot traffic starts 4 min into the test)

Another test was done with six home clients along with the two hotspot Wi-Fi clients just outside the house streaming HD YouTube videos. As in the previous test, the home hotspot clients started the streaming of the movies approximately 4 minutes after test began. This time, in an attempt to preempt the impact of the outside Hotspot clients, WMM was enabled and the home traffic was assigned to the Video Access Category thus giving the home clients the highest priority. Please note, the YouTube traffic towards the Hotspot clients was not manipulated by the access point. In other words any DSCP marking used by the origin server was carried over and mapped into the relevant WMM queue.

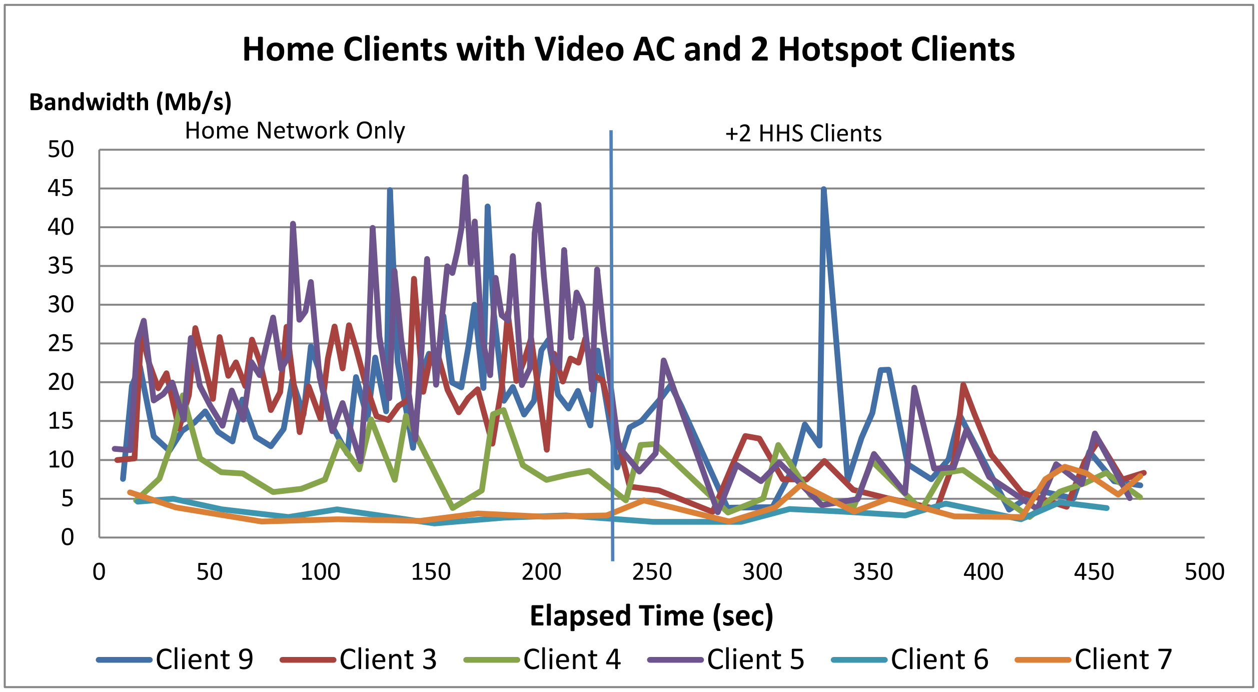

Figure 14: 6 home clients bandwidth
(Hotspot traffic starts 4 min into the test)

The impact of the outside clients was very noticeable, even with WMM applied. As seen by the chart below, the bandwidth of the home clients was reduced when the hotspot clients were active. Surprisingly, for some clients, the use of WMM priority actually resulted in lower bandwidth when compared to the client bandwidth without the use of WMM.

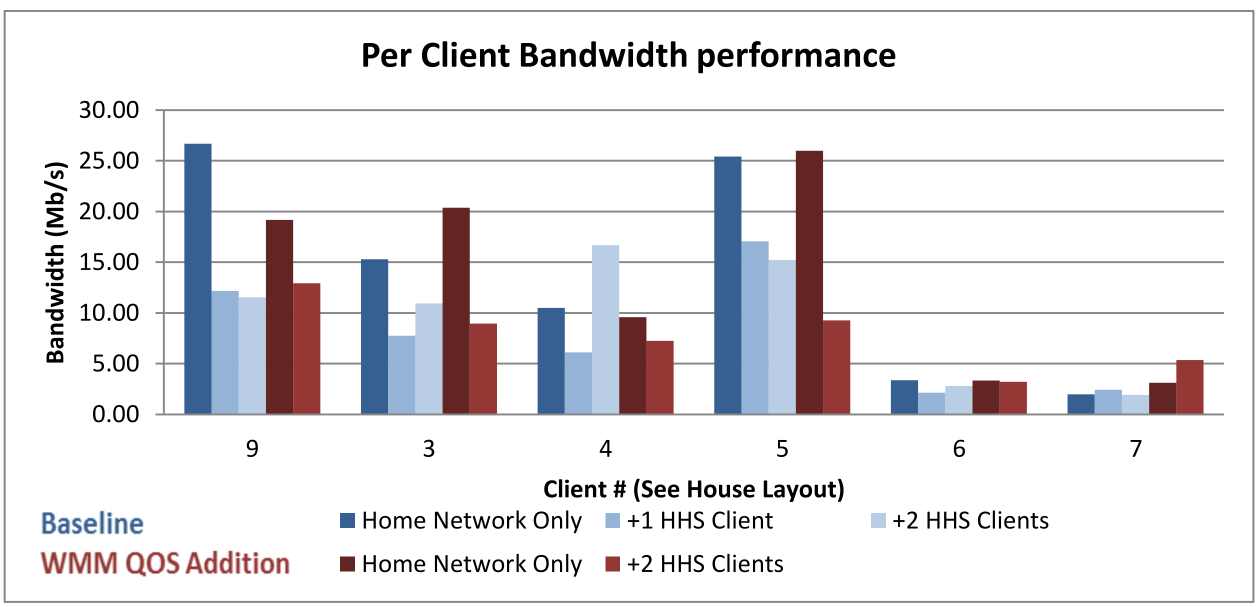

Figure 15: Per client bandwidth when hotspot clients are active with and without WMM

Bandwidth (Mb/s)					
Client	Baseline (no WMM)			With WMM QOS	
	Home Network Only	+1 HHS Client	+2 HHS Clients	Home Network Only	+2 HHS Clients
9	26.69	12.15	11.55	19.19	12.93
3	15.29	7.75	10.94	20.37	8.95
4	10.48	6.10	16.68	9.57	7.24
5	25.43	17.05	15.21	26.01	9.27
6	3.36	2.14	2.81	3.34	3.21
7	1.96	2.40	1.92	3.09	5.35
Aggregate	83.21	47.60	59.10	81.58	46.95

Figure 16: Comparison of per client bandwidth
when hotspot clients are active with and without WMM

Response Time (sec)					
Client	Baseline (no WMM)			With WMM QOS	
	Home Network Only	+1 HHS Client	+2 HHS Clients	Home Network Only	+2 HHS Clients
9	3.44	9.55	10.00	4.70	8.88
3	5.85	14.77	12.96	4.28	11.08
4	8.63	16.80	10.51	9.93	13.68
5	3.58	7.84	10.03	3.51	11.01
6	26.00	41.23	35.14	26.36	27.19
7	42.27	39.62	44.63	28.68	18.57
Average	14.96	21.64	20.54	12.91	15.07

Figure 17: Comparison of per client response time
when hotspot clients are active with and without WMM

As seen by this example a simple WMM mechanism that is based on extended backoff times has little positive impact on the bandwidth available to the home clients. The impact of remote clients with low MCS settings cannot be overcome by simply applying WMM. Customers experiencing such in-home Wi-Fi performance degradation will most likely have poor quality of experience, which may result in service calls to the service provider to “fix the Wi-Fi problem”.

The Impact of Home Security

A modern home security service usually involves the use of cameras that are located at the edges of the subscriber home and connected wirelessly to the home Access Point. Such devices and services are becoming more common than ever as found by our survey.(2) Other home security sensors may also be deployed, and most of them will probably be wirelessly connected. Although such cameras and sensors are always on, they usually send very low bit rate pictures or other information to a monitoring portal. At any time, however, the user may choose to view a live streaming feed from one or more of the connected cameras, which will require much more bandwidth to deliver the video. Security services therefore demand a changing level of allocated bandwidth in order to fulfill the promise of service. They also may require preferential treatment of these specific clients at the expense of other clients on the same in-home network. One of the key challenges here is the fact that the traffic from these IP cameras and sensors is upstream towards the access point. Any access point downstream WMM or other fairness mechanism does not apply, since the clients transmit upstream towards the access point based on the collisions and backoff mechanism inherent to Wi-Fi.

SOLVING THE IN-HOME WIRELSS QUALITY OF EXPERIENCE: MANAGING AIR TIME

As we have demonstrated in this paper, remote clients can impact a home network's overall performance severely. Lack of a proper quality of service monitoring mechanism can degrade the user experience and create severe subscriber retention problems. Standard Wi-Fi tools, such as WMM Access Categories, cannot mitigate the impact these clients have on network performance. A more sophisticated scheme that takes each client's time consumption of the shared air interface into consideration is needed. To provide the expected level of service, service providers must apply specific scheduling algorithms coupled with higher level application logic.

In order to minimize the impact of roaming hotspot clients, it is important to limit not only the number of such associated devices, but to also cap the amount of time they are allowed to use the air interface. One way of controlling Air Time is by allocating time based on SSID. For example, an AP might allocate only 10% of the total Air Time to the clients associated with the home hotspot SSID, while clients attached to the home SSID can enjoy 90% of the Air Time. The differentiation based on SSID, and the type of service that is associated with each SSID, is especially crucial when the service provider wants to deliver high-definition video over Wi-Fi.

In addition, an admission control mechanism for allowing or blocking slow clients (based on low bit-rate MCS) connection to the hotspot SSID is needed. Service providers can incorporate such mechanisms as part of the handshake with the associating client (e.g., RADIUS authentication of the requesting client). The AP is aware of its resources in terms of overall bandwidth and number of active clients. When the wireless media is congested, an AP may decide to disassociate or block a hotspot client from gaining access to the air interface.

In the case of Home Security, where specific, known clients need extra bandwidth, the AP should be configured to use the air interface in a manner that meets their minimum needed bandwidth. These clients may need preferred treatment by the Air Time scheduler depending on the overall network usage at that time. A pre-defined minimum bandwidth configuration may be needed to ensure the video streams these clients send arrive intact and with minimum delay. Different services require different solutions and logic; a dynamic Air Time management scheme should allow the service provider to allocate the Air Time resource between types of services associated with different SSIDs while differentiating between clients associated with the same SSID.

References

(1) Wi-Fi Multimedia™ Technical Specification Version 1.0 Wi-Fi Alliance

(2) Managing Cable TV Migration to IP – Part 1 Series Introduction: Market Drivers and Tech Challenges, S. Howe C. Cheevers, ARRIS