white paper

Optimizing Network and Client Performance Through Dynamic Airtime Scheduling

Table of Contents

Introduction

In wireless LANs one thing is certain, wireless performance is often not as advertised. Wireless is a shared medium, meaning that all clients and neighboring APs compete for the same limited bandwidth, in addition, each client's speed varies depending on the protocol it is running (802.11 a/b/g/n) and the signal strength, interference and noise it is experiencing. Older clients using lower speed protocols, interference, inconsistent RF coverage and clients connecting at the fringe of the network or moving behind obstructions all lead to low data rate connections. These slow clients consume more airtime to transfer a given amount of data, leaving less airtime for other clients, decreasing network capacity and significantly degrading the performance of all clients on the network. This paper reviews key issues that affect wireless LAN performance, and shows how a new patent pending wireless Quality of Service (QoS) technology from Aerohive Networks – **Dynamic Airtime Scheduling**— can solve these problems.

The benefits of Dynamic Airtime Scheduling are compelling to both the IT organization and to the users of the wireless LAN. It enables clients connected at higher data rates in a mixed data rate environment to achieve up to 10 times more throughput than they would get with traditional wireless LAN infrastructures - without penalizing lower speed clients. This means that the users see faster download times and improved application performance. It also means that low speed clients don't destroy the performance of the WLAN for the rest of the users. This allows IT to implement a phased upgrade to 802.11n and immediately start reaping the benefits of the new 802.11n infrastructure even if it takes years to upgrade all of the clients. And, because a user connecting at the fringe of the WLAN can no longer consume all of the airtime, the network impact of a bad client or a weak coverage area is diminished –allowing IT to reduce infrastructure investment, save IT time and increase user satisfaction.

Dynamic Airtime Scheduling, coupled with user and application based policy enforcement, allows IT to manage wireless network resources to transform a shared WLAN into a true multi-service network infrastructure that makes it possible to reliably move users and applications to the air.

Mixed Data Rates in Traditional Wireless LANs

The 802.11 standards allow for all wireless devices within range and on the same channel to compete equally for the wireless medium, allowing only one AP or client to communicate at a time. Once an AP or a client starts to transmit a wireless frame, all other wireless devices on the same channel must wait until the transmission is finished before they can transmit. After a wireless frame has been successfully transmitted and the receiving device sends the wireless ACK, all devices, including the one involved in the previous transmission, have equal opportunity to gain access to the channel and use it for transmission. If a device is transmitting, the period of time that another device needs to wait before trying to transmit is determined by the size of the frame being transmitted and the transmit and receive data rates between the client and its AP. For example, a wireless frame transmitted to or from a client connected at a low data rate may utilize 10 milliseconds of airtime, whereas it may take only 100 microseconds for a client connected at a high data rate. Even though the high speed client could have sent 100 frames in the time the slow client takes to send one frame, the fast client still has to compete fairly for the airtime on a frame by frame basis, so it spends most of its time sitting idly waiting for the slow client to finish so it can have another chance to transmit. Unfortunately this means that a single low speed client can slow down all of the other clients on the WLAN.

The following diagram (Diagram 1) demonstrates this point. It displays the elapsed time required to transmit eight frames to a high speed client, and at the same time transmit eight equal size frames to a lower speed client. Though the time required to transmit eight frames at a higher data rate is much shorter, the high speed client's frames cannot be transmitted while the lower speed client's frames are on the air.

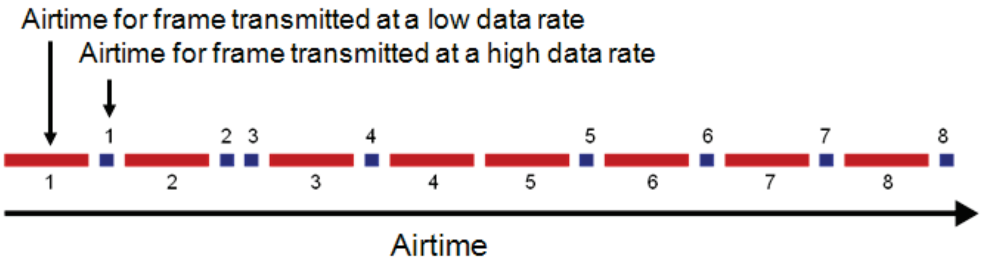

Diagram 1. **Frames over time based on standard wi-fi behavior**

The frames transmitted at the lower data rate take more time, but in the end, both clients receive the same number of frames, finish at approximately the same time and achieve the same throughput. This is, by no means, an equal use of airtime. The traffic to the lower speed client consumes much more airtime than the faster client and prevents the fast client from benefiting from its higher data rate.

If instead of giving an equal chance of transmitting a frame, you give equal airtime to clients, regardless of their data rate, the outcome can be improved for the higher speed client, with little to no impact on the lower speed client. The following diagram (Diagram 2) shows that by giving equal time slices to each client there can be many more frames transmitted to the higher speed client in the time it takes to transmit a single lower speed frame. In this example the higher speed client receives 4 frames for every frame sent to the lower speed client.

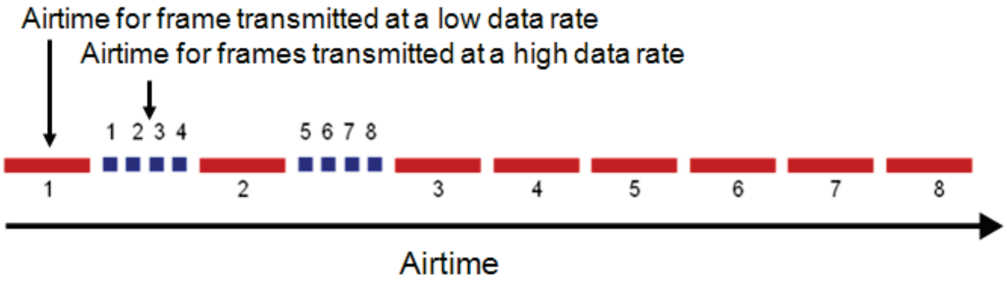

Diagram 2. **Frames over time based on equal airtime**

Over time, if both clients are downloading the same file, which in this example takes 8 frames to download, the higher speed client will finish much more quickly, leaving the rest of the airtime for the lower speed client. The low speed client takes approximately the same amount of time as it did in the previous example. If you factor in the added advantage of reduced contention, then the low speed client is also likely to finish more quickly. (Note: If you have more contention, the probability of collisions, random backoff times, and retransmissions increases thus lowering the performance for all clients that contend with each other. Therefore, getting the fast client off the air minimizes contention and helps to increase wireless LAN performance.)

This is demonstrated by the following tests (Diagram 3) conducted using a Veriwave™ WLAN test tool running a WiMix test. In the first test, the Veriwave connects a single wireless client and simulates a TCP-based HTTP file transfer from an Ethernet connected server by downloading 10,000 frames at 1500 bytes each. In the first test, the 802.11a wireless client is connected at 54Mbps, and in the second, 6Mbps. The results are shown in the following graphs.

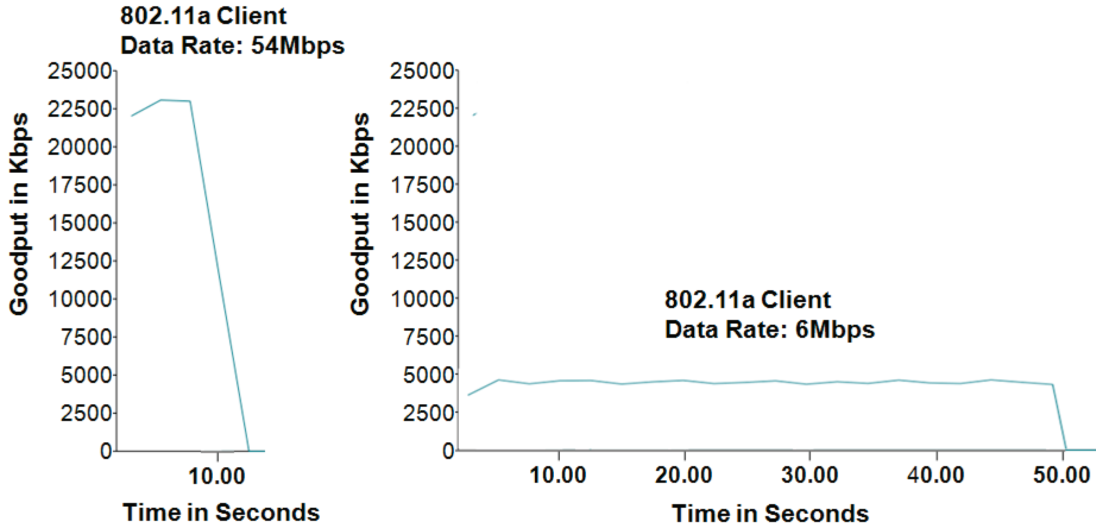

***Diagram 3*. Two simulated file transfers, one from a client connected at 54Mbps, and the other from a client connected at 6Mbps**

The 10,000 HTTP frames transmitted to the 54Mbps client takes approximately 12 seconds, while the transmission to the 6Mbps client with the same test takes approximately 50 seconds.

The next test (Diagram 4) shows what happens when the clients share the air while 10,000 frames are being transmitted to each client. The WiMix is setup so that the clients share the wireless airtime but there is no wireless contention.

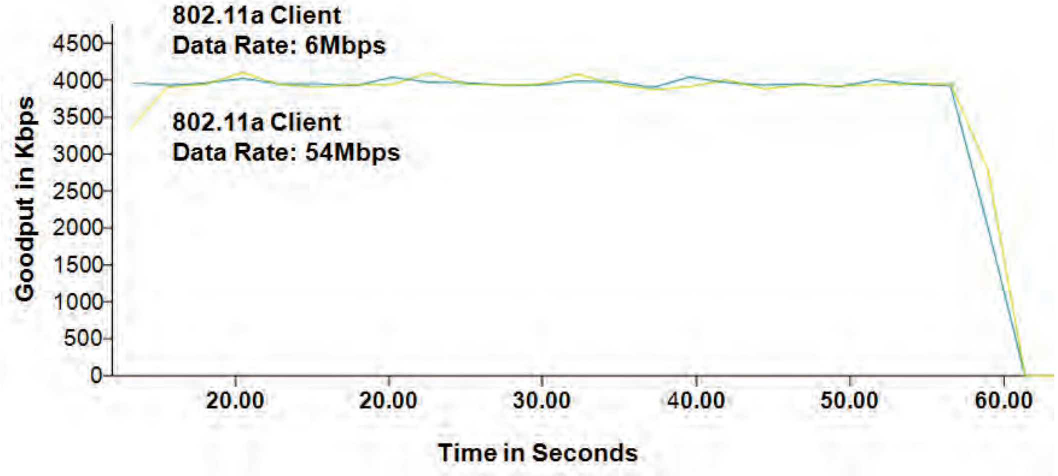

***Diagram 4.* Comparison of simultaneous transmissions of 10,000 1,500 byte HTTP frames to simulate a file downloaded by a low speed client and a high speed client**

As you would expect, the total time to complete the two downloads is approximately 62 seconds, the sum of the 12 and 50 second download times for the two individual transmissions. You can also see that the two clients had almost identical performance, both showing 4000Kbps (4Mbps) of goodput with each client taking approximately 62 seconds to complete the task. The lower speed client consumed the airtime and slowed the faster client down to its speed, leading to very disappointing performance for the higher speed client.

This problem can be dramatically improved by granting each client an equal amount of airtime rather than an equal chance of transmitting a packet. This minimizes the ability for low speed clients to reduce the performance of higher speed clients, while still giving low speed clients an equal share of airtime. Aerohive's patent-pending Dynamic Airtime Scheduling accomplishes this goal by scheduling airtime based on IT-specified policies and improves client and network performance. The following chart (Diagram 5) shows the same two-client test case, but now with Dynamic Airtime Scheduling enabled.

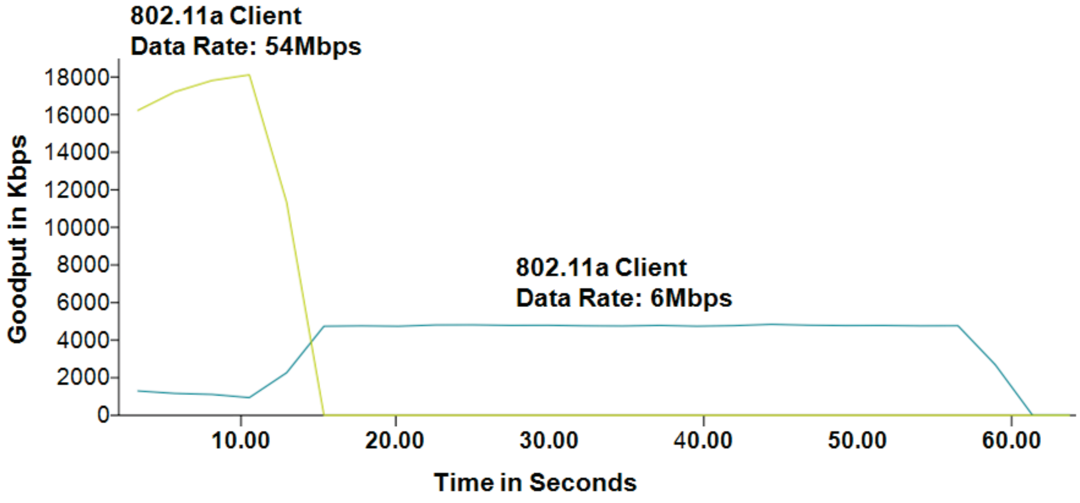

***Diagram 5.* Comparison of simultaneous transmissions of 10,000 1,500 byte HTTP frames to simulate a file downloaded by a low speed client and a high speed client with Aerohive's Dynamic Airtime Scheduling**

The test shows that by scheduling airtime, the higher speed client finishes 4 times faster – a 300% performance increase, while the lower speed client finishes at approximately the same time. Fast client performance is dramatically improved and overall network capacity is increased without penalizing low speed clients.

Aerohive QoS

Aerohive's Dynamic Airtime Scheduling is built upon an existing capable and flexible QoS engine. The Aerohive Quality of Service (QoS) engines within HiveAPs provide highly granular prioritization and deterministic transmission of packets onto the wireless and wired networks. With Dynamic Airtime Scheduling, QoS is used to improve performance, but it also serves the other purpose of ensuring critical applications, such as voice, are handled with expediency. The Aerohive QoS capability consists of five main components:

1. **Classifier and Marker** –categorizes traffic into eight queues per user based on QoS classification policies. These policies can be configured to map traffic to queues based on network service, MAC OUI (Organization Unique Identifier), SSID and interface, or priority markings on incoming packets using IEEE 802.1p, 802.11e or DiffServ. The classifier is also responsible for marking traffic with IEEE 802.11e or DiffServ so it can be prioritized through the wireless LAN. Traffic going out an Ethernet interface can be marked with IEEE 802.1p or DiffServ.

2. **Policer** – rate limits traffic on a per user, per user queue or per group basis to prevent a user, or an entire class of users (e.g., guests) from consuming excess network resources.

3. **User Queues** –8 queues per user to allow for granular prioritization of traffic, and weighted access among clients

4. **Scheduler** – uses strict priority and weighted round robin techniques to granularly schedule traffic from each of the eight queues into the Wireless Multi-Media (WMM) hardware queues. It does this by taking into account the configured weight of the user profile that is assigned to the user, and the configured weight of each of the user's eight queues. Because the QoS packet scheduling engine runs on the HiveAP (rather than on a remote WLAN controller), it has the ability to closely monitor the availability of the WMM queues and instantly react to changing network conditions. The QoS packet scheduling engine only transmits to WMM queues when they are available, queuing packets in eight queues per user in the meantime. This prevents dropped packets and jitter, which adversely affect time-sensitive applications such as voice, and prevents TCP flow performance degradation caused by contention window back-off algorithms, which is when TCP packets are dropped.

5. **Wireless Multi-Media Queues (WMM)** –a standard mechanism for transmitting packets from hardware queues prioritized by four access categories: Voice (highest priority), Video, Best Effort, and Back Ground (lowest priority), to the wireless medium. Packets

from higher priority access categories are transmitted with a smaller inter-frame space and a random back-off window to allow transmission to the wireless medium with less delay.

These QoS engines are designed to add intelligence and control to a shared wireless network so that it can become a true multi-service network infrastructure capable of supporting a diverse range of users and applications. They allow bandwidth to be allocated to different users and groups – for example ensuring guests get 1/10 of the bandwidth that employees get during times of congestion but providing more bandwidth during quiescent times. And they ensure that time and latency sensitive traffic, such as voice and video, has priority over other types of traffic –e.g. file transfers and email— that are not as sensitive. These mechanisms can be implemented in a wireless LAN to help to assure that the voice quality for their VoWLAN phones or the quality of their video over the wireless LAN will not be impaired by large data transfers.

These bandwidth based QoS mechanisms are very effective when clients are running at similar Wi-Fi data rates, or when airtime is not the bottleneck. However, clients at different data rates that consume different amounts of airtime, and slow clients consuming all the airtime and impairing network performance means that these bandwidth based QoS mechanisms are not sufficient to transform a shared WLAN into a true multi-service network infrastructure. Airtime needs to be managed and scheduled to truly enable a reliable, high performance, multi-service WLAN. Aerohive has made this a reality by delivering Dynamic Airtime Scheduling, a major enhancement to the QoS scheduling engine that allows the QoS scheduling engine to react based upon airtime consumption rather than bandwidth consumption.

The QoS engines allow airtime scheduling to be used in conjunction with the rest of the HiveAP's QoS capabilities, so that for example, critical voice traffic can still be treated with strict priority and forwarded ahead of other traffic, even if that voice client is using more airtime than it should. Combined with a flexible policy engine, IT has tremendous ability to implement user, group, device and application-based network bandwidth and airtime management policies.

Dynamic Airtime Scheduling Concepts and Examples

Using bandwidth-based QoS scheduling mechanisms, or with no QoS at all, if traffic is being transmitted to or from clients connected at differing data rates, the faster clients throughput will slow down to the rate of the lowest speed client. This happens whether all clients are of the same type or if there is a mix of 802.11a/b/g/n clients. To demonstrate this, we used Veriwave WiMix to conduct tests to show the results of situations with and without Dynamic Airtime Scheduling enabled.

Single Protocol (802.11a) with Different Data Rate Clients

In the first test (Diagram 6), we connect 3 clients that are all running the same 802.11a Wi-Fi protocol, but at three different data rates - 54Mbps, 12Mbps and 6Mbps. This simulates clients connected to the same AP but at varying distances or interference levels which cause the data rate to change. These tests simulate transferring 10,000 good 1500 byte HTTP packets for each client, showing the TCP goodput (traffic throughput that reaches its destination intact). The graph on the left shows that even though two of the clients are transmitting at a higher data rate, without airtime scheduling, they end up with throughput equal to the low data rate client. It takes all three clients 88 seconds to finish transferring 10,000 HTTP.

Note: This test is conducted in a closed environment, and the clients wait their turn without contention. In a real life open air environment, the client contention would cause longer transfer times, though the behavior remains the same.

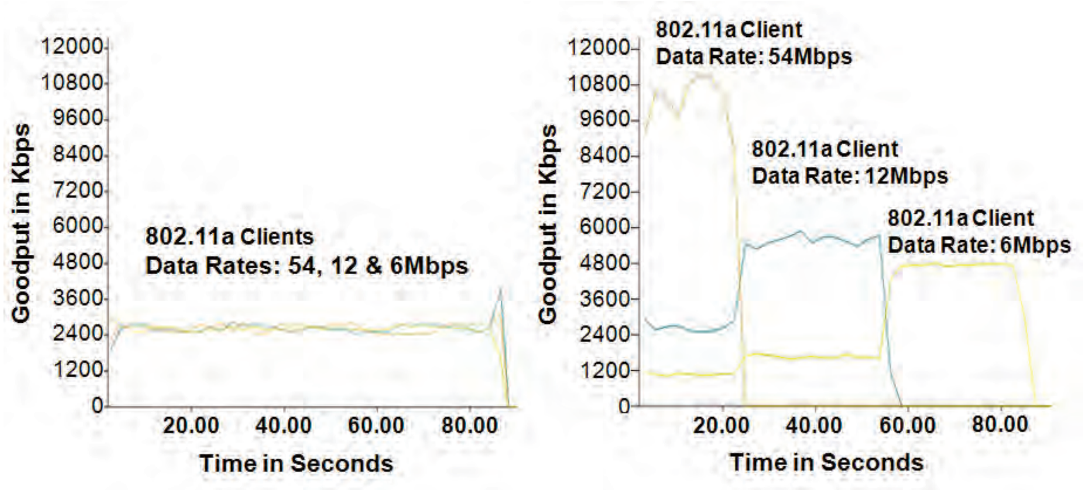

***Diagram 6.* Three simultaneous file transfers, one test without and one test with Dynamic Airtime Scheduling**

The graph on the right shows the same test with Aerohive's Dynamic Airtime Scheduling enabled. You can see that the transfer time for the 54Mbps data rate client is 4 times faster dropping from 88 seconds to 22 seconds. It finishes its transfer and frees up additional airtime for the remaining two clients to use. The 12Mbps client can now transmit more frequently, allowing it to finish in 2/3s the time at 59 seconds. Then the 6Mbps client can finish its transfer without contending for airtime, so it can transmit more quickly and it finishes at exactly the same time as in the first test.

802.11n Environments

The next test (Diagram 7) adds 802.11n clients to the equation to simulate a typical WLAN's transition to 802.11n – where 802.11n APs are servicing a mixture of 802.11n clients and legacy (a/b/g) clients. These tests show how Dynamic Airtime Scheduling enables a high performance 802.11n wireless LAN to not be hampered by clients connected at lower data rates – even if some of those slow clients are slow 802.11n clients. Depending on distance to the AP, position of the antenna or interference, the data rate of a connected .11n client can range from 13.5 to 270Mbps, so some 802.11n clients will be slower than other 802.11n client, and it is even possible that some 802.11n clients could actually be slower than fast a or g clients.

To simulate this mixed n/a environment, we connect three 802.11n clients at 3 different rates on the 5GHz radio band of an AP. The rates used are 270Mbps, 108Mbps and 54Mbps (with frame aggregation enabled). Simultaneously we connect three 802.11a clients at 54Mbps, 12Mbps and 6Mbps. This simulates six mixed clients connected to the same AP but at varying distances or interference levels which cause the data rate to change. These tests transfer 10,000 good 1500 byte HTTP packets for each client. The graph to the left shows the results without airtime scheduling. As we saw in the previous

802.11a tests, each of the clients, though connected at different data rates, drops down to the throughput of the client at the lowest rate, which in this case is 6Mbps. The time it takes all 6 clients to finish transferring 10,000 HTTP packets is between 90 and 110 seconds.

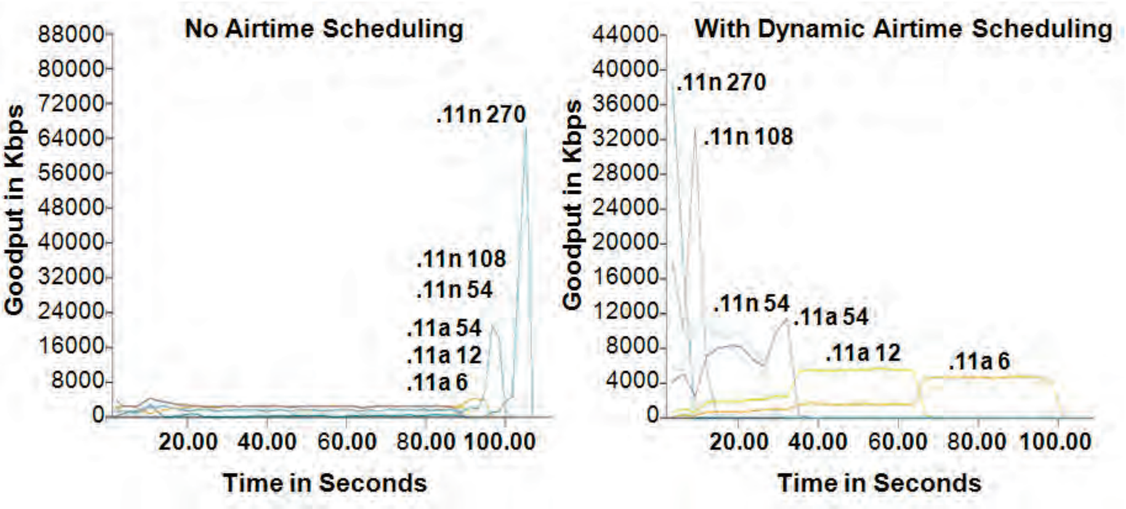

Diagram 7. **Dynamic Airtime Scheduling in 802.11 environments**

The graph to the right shows the same test with a HiveAP using Aerohive's Dynamic Airtime Scheduling. The transfer time at the 270Mbps data rate is approximately 10 seconds – about 10 times faster than the 110 seconds seen in the previous test. Likewise, the rest of the transfer times improved significantly. The .11n(108Mbps) transfer was over 6 times faster, the .11n at(54Mbps) transfer was over 3x faster, the .11a(54Mbps) transfer was 2.5 times faster, the .11a(12Mbps) transfer was 30% faster, and the .11a(6Mbps) transfer decreased slightly (10%).

As with the first test, with Dynamic Airtime Scheduling network performance is dramatically improved. All of the higher data rate clients saw substantial improvements in performance, while the low speed client saw almost no negative impact. In an open air network, the effects of the performance gain will even be higher, because once the higher speed clients finish, fewer clients are on the air, contention and retries decrease leading to performance increases.

Dynamic Airtime Scheduling in a Nutshell

Now that you have seen what Dynamic Airtime Scheduling can do, let's run through the main concepts of how it works. With bandwidth-based scheduling, the AP calculates the bandwidth used by clients based on the size and number of frames transmitted to or from a client. Bandwidth-based scheduling does not take into account the time it takes for a frame to be transmitted over the air. As discussed in previous sections, clients connected at different data rates take different amounts of airtime to transmit the same amount of data.

By enabling Dynamic Airtime Scheduling, the scheduler allocates airtime instead of bandwidth to each type of user, user, and user queue, which can be given weighted preferences based on QoS policy settings. When traffic is transmitted to or from a client, the HiveAP calculates the airtime utilization based on intricate knowledge of the clients, user queues, per packet client data rates, and frame transmission times, and ensures that the appropriate amount of airtime is provided to clients based on their QoS policy settings.

Dynamic Airtime Scheduling is made possible because it is performed directly on the HiveAPs responsible for processing the wireless frames. This gives the scheduling engine access to all the information needed in real time (at the microsecond level)), allowing the HiveAP to react to instantaneous changes in client airtime utilization, that occur when the client is moving or even stationary.

Upstream Traffic and Dynamic Airtime Scheduling

While the Wi-Fi standard does not currently allow an access point to force a client to not transmit, Dynamic Airtime Scheduling is still able to schedule and control upstream traffic. It measures the total airtime of the client including upstream traffic, and by aggregating both the sent and received traffic Dynamic Airtime Scheduling can ensure that client airtime consumption for both send and receive does not exceed its allotment. If a client is uploading a large file and attempting to consume more than its allotted airtime, the scheduling engine will queue that client's downstream traffic, which will delay the transmission of protocol ACKs and slow the client's transmission. The precision of scheduling is less than on the directly controlled downstream traffic, but the end result is that upstream airtime can be scheduled and controlled without having to resort to non-standard Wi-Fi approaches which can create client interoperability problems and interfere with neighboring networks.

Voice over WLAN with Dynamic Airtime Scheduling

When voice over WLAN services are used, Dynamic Airtime Scheduling is a complimentary technology that reduces contention in a mixed data rate environment, but it is not directly utilized for voice traffic. HiveAPs classify voice traffic into queues that use strict priority instead of scheduling with weighted round robin. The airtime for strict priority traffic is measured by the scheduling engine so that it can appropriately deduct the airtime used by voice in its scheduling decisions, but it does not schedule the transmissions of voice packets, instead it sends them directly to the WMM voice queue for immediate transmission. Because voice traffic uses small packets and is low bandwidth it does not consume much airtime and should only be transmitted with strict priority.

To demonstrate how voice calls are handled mixed in with clients using dynamic airtime scheduling, Veriwave WiMix is used to simulate 20 bidirectional VoIP calls using SIP (Session Initiation Protocol), while six clients connected at mixed 802.11n data rates are simultaneously downloading files using HTTP. The following graph (Diagram 8) shows six clients downloading 20,000 HTTP packets with two clients connected at 270Mbps, two clients connected at 108Mbps, and two clients connected at 54Mbps. Dynamic airtime scheduling ensures that the higher speed transfers finish quickly, freeing up the air for transfers at lower speeds.

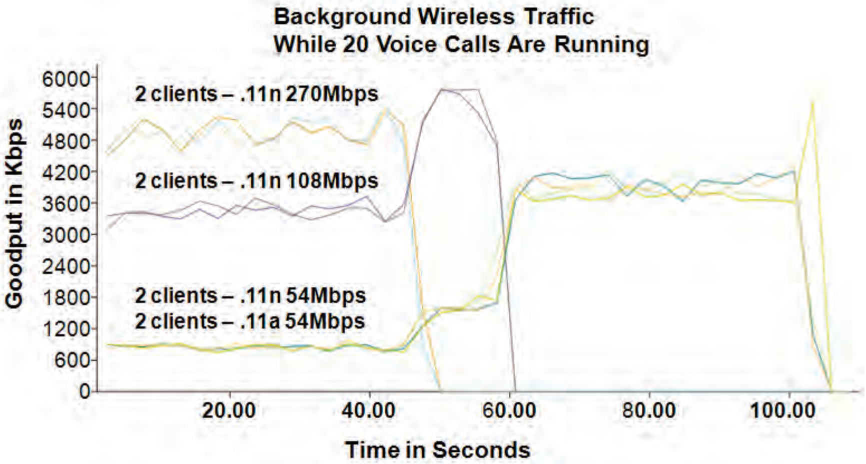

Diagram 8. **Eight clients downloading files connected at different data rates (shown) while twenty other clients are simultaneously on voice over WLAN calls with toll quality (not shown)**

Simultaneously, the test tool simulates 20 wireless voice clients connected at a 54Mbps data rate, with bidirectional SIP voice calls. The following graph (Diagram 9) of the test output shows that each client maintains a consistent 88Kbps rate required for exceptional voice quality.

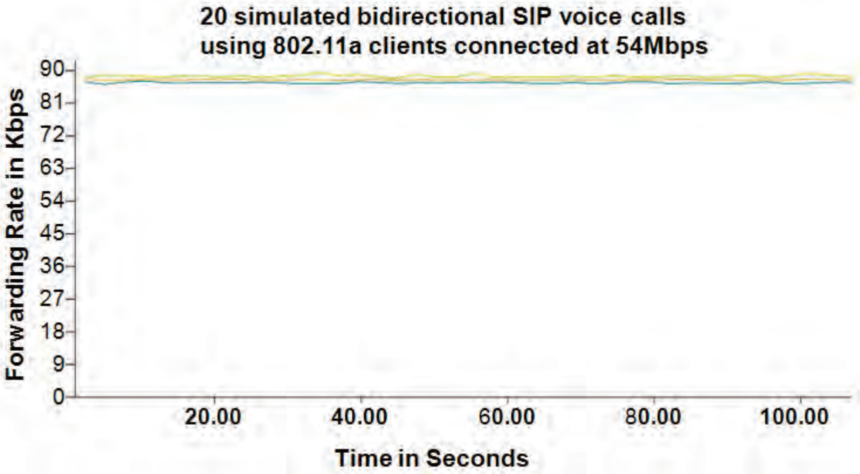

Diagram 9. **Example of twenty simulated SIP voice calls getting toll quality running concurrently with the eight simulated file transfers shown previously**

The test output from WiMix also shows that the Voice quality is exceptional with MOS scores for each call at 4.2 and R-Values at 85.5.

Giving Preference with Dynamic Airtime Scheduling

The use of dynamic airtime scheduling is not just reserved for ensuring clients have equal airtime, it can also be used to give preference to users, types of users, SSIDs, client devices and network services. This can be used to give employees more airtime than guests if there is contention between them. For an example, let's say you have two SSIDs, one for employees and one for guests. You can set the weight of the employee SSID higher than that of the guest SSID. If there is no wireless contention, both employees and guests get full use of the airtime. If employees and guests are using the wireless at the same time, the employees will get more airtime. The following picture (Diagram 10) shows the results of a Veriwave WiMix test that simulates simultaneous file downloads of 9 employees and 9 guests. To simplify the output of the test, employees and guests have 3 clients each connected at 270, 108, and 54Mbps data rates all downloading the same file.

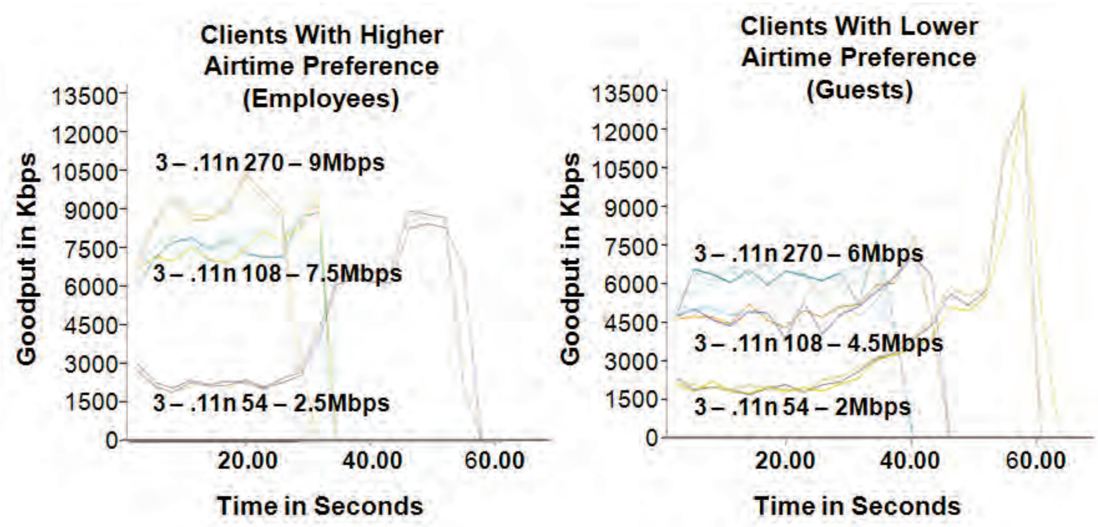

***Diagram 10.* Airtime preference given to employees vs. guests; Both the above graphs are from a single test run, it is shown as two separate graphs for simplicity given the number clients involved**

By giving the employee SSID more weight, you can see from the output that the employees get more airtime than guests, but within each group airtime scheduling still optimizes throughput of the clients running at mixed data rates. The weight preferences can be fine tuned to provide as much preference as desired.

True Per-Packet Airtime Calculations vs. Protocol Based Scheduling Using Ratios

With Dynamic Airtime Scheduling, the actual airtime in microseconds that is consumed by every frame is measured so that the scheduling engine makes its decisions based on true airtime consumption. Regardless of what wireless protocol is used, IEEE 802.11a/b/g/n, or what data rate clients are running at, each client is given its appropriate share of airtime based on actual airtime usage. This is significantly more accurate than Protocol Based Scheduling approaches that assume all clients running a given Wi-Fi protocol are running at the same data rate, and schedule traffic based on a fixed ratio between each Wi-Fi protocol.

If a QoS system uses this kind of simplistic Protocol Based Scheduling that blindly assumes that .11g traffic is faster than.11b traffic, or simply prioritizes .11n traffic over .11g traffic, and does not take into account the actual data rate for the connected clients, the performance of the network can be adversely affected. In a worst case scenario, Protocol-based Scheduling can actually make a network perform slower than with no

QoS at all! In addition, Protocol-based scheduling does nothing to address the issue of slow 802.11n clients impacting faster 802.11n clients.

The following picture (Diagram 11) shows that Protocol-based Scheduling does nothing to address fast and slow clients running the same Wi-Fi protocol (e.g. 802.11n). The graph shows file downloads from two clients - a slower 802.11n client connected at a 54Mbps data rate, and a fast 802.11n client connected at a 270Mbps data rate. Here the lower rate .11n client is able to slow the faster client, just as discussed earlier in the paper when no airtime scheduling was enabled. This issue becomes particularly important with the transition to .11n where more and more clients deployed in the network will be .11n and where the data between a fast and slow client can vary so substantially (e.g. from 270Mbps to 13.5Mbps).

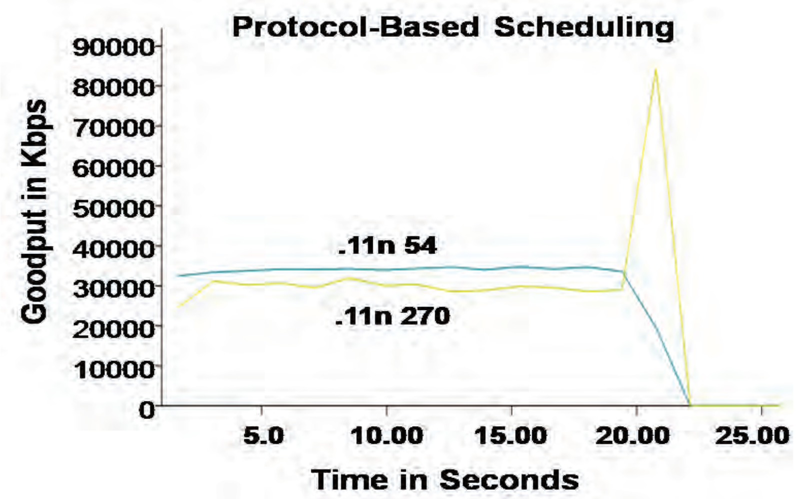

***Diagram 11.* Protocol-based scheduling with different speed clients of the same protocol**

True per-packet airtime scheduling will always optimize airtime utilization no matter the environment, while simplistic protocol-based scheduling systems must make coarse assumptions which limit the applicability, and performance benefit of the feature.

Conclusion

Aerohive's Dynamic Airtime Scheduling is an exciting and innovative new technology that provides QoS based on airtime instead of just bandwidth. Managing airtime is critical because airtime consumption affects all of the clients on the network. It needs to be dynamically managed because it varies widely across all of the clients on the network – not just because clients are running different Wi-Fi protocols, but also because of instantaneous changes in relative closeness to the AP, signal strength, interference and error rates. With Dynamic Airtime Scheduling, airtime can be dynamically scheduled to increase network and client performance and to allocate network capacity based on IT-specified policies.

Dynamic Airtime Scheduling increases the performance of wireless networks and transforms a wireless LAN into a true multiservice network infrastructure.

About Aerohive

Aerohive Networks reduces the cost and complexity of today's networks with cloud-enabled, distributed Wi-Fi and routing solutions for enterprises and medium sized companies including branch offices and teleworkers. Aerohive's award-winning cooperative control Wi-Fi architecture, public or private cloud-enabled network management, routing and VPN solutions eliminate costly controllers and single points of failure. This gives its customers mission critical reliability with granular security and policy enforcement and the ability to start small and expand without limitations. Aerohive was founded in 2006 and is headquartered in Sunnyvale, Calif. The company's investors include Kleiner Perkins Caufield & Byers, Lightspeed Venture Partners, Northern Light Venture Capital and New Enterprise Associates, Inc. (NEA).

Corporate Headquarters
Aerohive Networks, Inc.
330 Gibraltar Drive
Sunnyvale, California 94089 USA
Phone: 408.510.6100
Toll Free: 1.866.918.9918
Fax: 408.510.6199
info@aerohive.com
www.aerohive.com

EMEA Headquarters
Aerohive Networks Europe LTD
Sequel House
The Hart
Surrey, UK GU9 7HW
+44 (0)1252 736590
Fax: +44 (0)1252711901

WP1101103