Whitepaper

Samsung Wireless Enterprise

Samsung
Smart WLAN Solution

Smart Capacity & Security
for Smarter Mobility

AirEqualizer

THE NEW
BUSINESS
EXPERIENCE

SAMSUNG

Introduction

In today's world, it's more evident than ever that wireless communications have become one of the most popular means of personal and business communication. Wireless devices such as smartphones, tablets or laptops are used on a daily basis in a diverse range of scenarios from coffee shops, airports and homes to meeting rooms and office desks. In the enterprise world, communication goes beyond basic email and web browsing into a whole new set of applications such as VoIP, video conferencing and streaming. And more importantly; all of these services need to be available whilst roaming from location to location and without degradation or loss of service - in short: complete mobility.

The original wired Local Area Networks (LAN), which have been the mainstay in office workspaces for many years, have become obsolete and cannot cope with this rapidly growing group of applications that modern businesses require to succeed in the real world, since they cannot offer mobility and flexibility. Wireless LANs (WLAN) have taken over the prime role of delivering the communication platform and consequently, people rely on them more than ever.

Due to the dramatic increase in demand, WLANs are developing incredibly fast in order to minimize the performance gap or limitations they have compared to traditional wired networks. The source of these limitations lies in the shared medium nature of the communications. This causes interference, which can result in unreliability of the service. WLAN standards have evolved very fast, with a main goal of increasing speed and reliability - from the early 802.11 to b, g, a, n and the new 802.11ac.

It is in enterprise networks where different devices from different vendors and models coexist, sharing the same wireless resources. One of the design challenges of WLANs is the way channel access mechanisms schedule clients to access the medium. The access mechanisms have to take into account a "fairness" element when distributing wireless resources among clients. In other words, the network has to make sure that all the clients are treated equally.

However, even though the equality term might seem like a common-sense idea, the way in which fairness and equality are defined can drastically affect the behavior and consequently the performance of the entire network.

For example, the consequences of implementing fairness algorithms based on "equal chance to transmit" over a period of time are completely different from fairness algorithms based on "equal amount of data transferred" over a period of time, as it will be illustrated in this paper.

In this paper we summarize the role of the Samsung AirEqualizer algorithm in providing an efficient way of guaranteeing fairness in the enterprise network while maximizing the network efficiency. Initially, a brief description of the technical background is provided, followed by a summary of the AirEqualizer algorithm and its performance. Test examples are shown in the final sections to illustrate the values this algorithm has, and of course its competitive advantages.

Evolution of 802.11

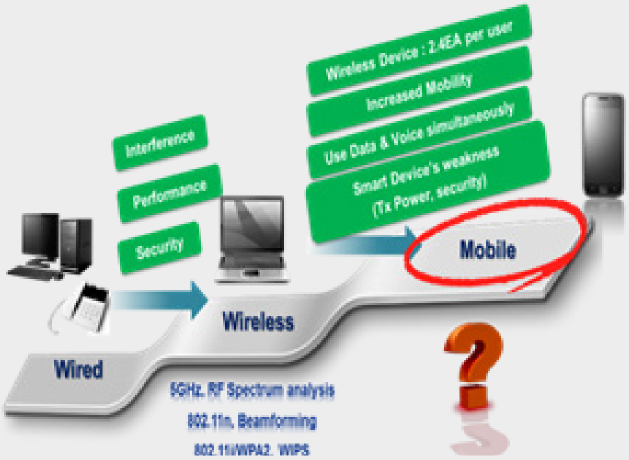

In the beginning of Internet communication services, enterprises installed wired LANs to allow clients to access the Internet through their fixed desktops. The versatility and flexibility of WLAN for enterprise environment were soon made obvious. Enterprise WLANs allow the use of data traffic and voice communication indoors, outdoors or a combination of both. However, the performance of wired LANs was far superior to WLANs. Both academia and industry made numerous efforts to reduce this gap to a minimum, and the factors they focused on were the following:

- Performance: The data speed (i.e., throughput) obtained with WLANs was considerably lower than the achievable throughput for wired networks. New standards with enhanced modulation and coding techniques were developed, increasing the speed and efficiency of the data transmissions for WLANs.

- Interference: WLAN technology shares the same radio band with Bluetooth, microwave ovens, cordless DECT phones, and even a wireless mouse. Sometimes, it even interferes with the lighting system. The abundance of interfering products in wireless radio ba nds was tackled by including different spectrum bands with less interfering occupation (i.e., 5GHz radio band) and also RF spectrum analyzers, which act against interference.

- Security: Security against malicious attacks is a big concern when using WLANs and news about Wi-Fi hacking was prominent. The 802.11i standards and Wireless Intrusion Prevention systems (WIPS) were developed to protect wireless communications.

In this paper we focus on the efforts to improve the communication speed of stations (Note that, throughout the paper, the terms client and station are used indistinctively). The following table provides a quick summary of the 802.11 evolution and how different transmission schemes successfully enhance the communication speed:

PHY	Transmission Schemes	Frequency Bands	Supported Transmission Rates (Mbps)
Baseline	DSSS, FHSS and IR	2.4 GHz (DSSS & FHSS)	1, 2
802.11a	OFDM	5 GHz	6, 9, 12, 18, 24, 36, 48, 54
802.11b	CCK	2.4 GHz	5.5, 11, DSSS rates
802.11g	OFDM		6, 9, 12, 18, 24, 36, 48, 54, 802.11b rates
802.11n	OFDM + MIMO	2.4 and 5 GHz	Up to 600
802.11ac	OFDM +DL MU-MIMO	5 GHz	Up to 6933.3

With all these improvements the performance gap between wired and wireless networks has been drastically reduced. However, not all of the stations in a company workspace use the same transmission schemes. The potential problems arising from the coexistence of newer or fast stations with slower or legacy stations should be considered to prevent unsatisfactory service to the clients or a critical degradation of the overall network performance.

Performance Anomaly

Enterprises are very heterogeneous environments, where a multitude of stations coexists under the same roof sharing the same wireless resources. These stations implement different WLAN standards, for example some may use the IEEE 802.11b (2.4 GHz) or IEEE 802.11n (5GHz). Depending on the location, stations use different modulation and coding techniques designed to increase the speed of the transmission under favorable medium conditions, increasing the communication efficiency while guaranteeing the reliability, so within the same enterprise network we can have fast and slow clients.

The explanation of this phenomenon lies in the original design of the 802.11 channel access mechanisms and how the station clients gain access to wireless resources. We define wireless resources as the period of time the network allocates for clients to transmit their data packets, also known as airtime resource. Since all the stations in the network have to share the medium and therefore can only transmit one at a time, the resource is the airtime that they consume in the transmission of their data packets when they access the channel. The IEEE 802.11 standard defines two main channel access mechanisms, a not-very popular centralized one called Point Coordination Function, where the AP allocates the airtime resources for each client and the Distributed Coordination Function that utilizes the basic Channel Sensing Multiple Access with Collision Avoidance commonly known as CSMA/CA principle to grant the access of a client to the medium. In order to support delay-sensitive traffic, the Enhanced Distributed Channel Access (EDCA) built on top of DFC was implemented, also based on CSMA/CA.

The EDCA transmission procedure can be summarized as following: once the scheduler has selected the next packet to transmit from one of the four Access Category (AC) queues, belonging to different types of traffic (Voice, Video, Best Effort and Background), the entire procedure commences. The transmitter MAC senses the state of the channel to ascertain whether it is in use or not. If the channel is detected busy, the transmitter MAC waits until it becomes idle. Once the transition to idle occurs, it postpones the transmission for Distributed Inter-Frame Spacing (DIFS). If the channel stays idle during this period, then the transmitter MAC initializes its back-off timer by choosing a random integer drawn from an uniform distribution over [0, CW], where CW is the Contention Window size ranging from CWmin to CWmax.

The back-off timer is decremented every slot time interval that the sensed medium remains idle. The packet transmission starts when the counter reaches zero. The function of the back-off timer is to reduce the chances of simultaneous transmission attempts by the contending stations which would result in collisions and consequently transmission failure and waste of airtime resources.

After each successful reception of a packet, the receiver acknowledges the transmission with a reply message known as (ACK) frame after a Short Inter-Frame Spacing (SIFS). If the packet is not received correctly, then the ACK frame is not transmitted and the transmitter increases the Contention Window size reducing the collision probability. This channel access mechanism gives all the stations the same probability to access the channel, which could be considered "fair" approach in terms of fair access to the wireless medium.

However, the amount of airtime resources required to successfully transmit the same amount of data is heavily dependent on the PHY standard and the modulation and coding scheme (MCS) used by the station. The transmission scheme used by the station determines the speed at which the data bits can be transmitted from the transmitter to the receiver, which translates into different amounts of airtime resource consumption to transmit the same amount of data. We can study the effect of the MAC combined with the different PHY standard implementations in the enterprise WLAN behavior with different examples:

- Fast stations coexist with slow stations.

  In these heterogeneous mobile wireless environments, a very common phenomenon known as "performance anomaly" can be observed where the overall expected performance of the network is seriously degraded. In the worst case scenarios, where some stations are located near the AP and able to use higher data rates while some other stations are located near the network boundaries and forced to use low transmission rates to assure reliability, one station that uses a lower bit rate competes with the other stations and the throughput of all stations may be significantly limited. One example would be an office with laptops near the AP and mobile smartphones located near the AP's coverage boundary. Here, the fast stations see their throughput decreased drastically due to the presence of slow stations.

As an illustrative explanation you can observe in the table below the amount of airtime resources required to transmit a 1000 bytes TCP data frame depending on the standard and transmission scheme used:

IEEE 802.11b

PHY rate	1.0 Mbps	2.0 Mbps	5.5 Mbps	11.0 Mbps
TX time (us)	8946.5	4578.5	1869.5	1129.5

IEEE 802.11a

PHY rate	6.0 Mbps	9.0 Mbps	12.0 Mbps	18.0 Mbps	24.0 Mbps	26.0 Mbps	48.0 Mbps	52.0 Mbps
TX time (us)	1602.5	1130.5	886.5	650.5	530.5	410.5	354.5	334.5

IEEE 802.11n, 2.4GHz

PHY rate	6.05 Mbps	13.0 Mbps	18.8 Mbps	26.0 Mbps	24.0 Mbps	39.0 Mbps	48.0 Mbps	52.0 Mbps
TX time (us)	1510	846.5	630.5	518.5	410.5	354.5	338.5	322.5

IEEE 802.11n, 5GHz

PHY rate	6.5 Mbps	9.0 Mbps	12.0 Mbps	18.0 Mbps	24.0 Mbps	26.0 Mbps	48.0 Mbps	52.0 Mbps
TX time (us)	1510	846.5	630.5	518.5	410.5	354.5	338.5	322.5

This hopefully makes obvious the benefit of the standard improvement and the availability of several transmission rates, since with higher transmission rates, larger number of data frames can be transmitted under the same amount of time. However, as it has been explained before, the original CSMA/CA does not take this into account and gives the same access probability to each station in the network regardless of its transmission capabilities. Therefore slow stations consume the airtime resources for a considerably longer period of time to transmit the same amount of data, preventing the fast stations from taking advantage of their faster speed and eliminating the advantages of using higher rates. The throughput achieved by fast stations is bounded by the slow ones, which degrades the overall performance perceived by the stations in the WLAN.

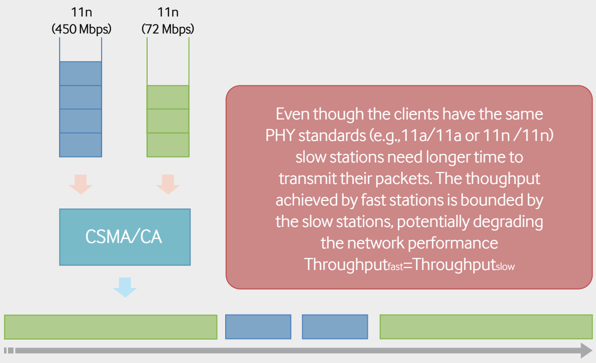

- IEEE 802.11n stations share medium with legacy stations.
  Before the aggregated-MPDU (AMPDU) procedure used in 802.11n, the effect of the performance anomaly phenomenon was of the order of the slow station throughput. However, AMPDU reduced the impact considerably by allowing the transmission of multiple aggregated frames at once in a super frame, but the problem was not solved. When fast stations aggregate N packets their throughput is N times higher than the slow stations. However, the slow transmission can still be longer than the aggregated packed transmission so the throughput is still bounded by the slow stations. This behavior of slow stations still penalizes fast stations and privileges the slow ones, decreasing the overall network speed.

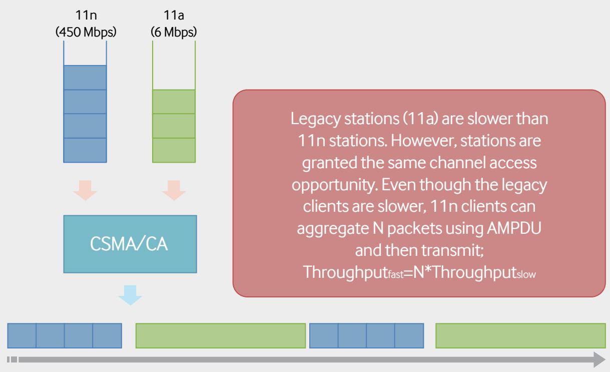

Fairness

It might seem odd to question the term fairness, since its benefits would seem rather evident, but the following series of examples are going to illustrate the consequences of different channel access mechanism implementations in the WLAN. We can observe in the diagram below, examples of homogeneous enterprise WLANs and the theoretical goodput values obtained when one or two stations are present. In homogeneous scenarios, all stations are treated equally by the AP and therefore fairness is easily achieved whether stations are all slow (e.g., 150 Mbps) or all fast (e.g., 450 Mbps).

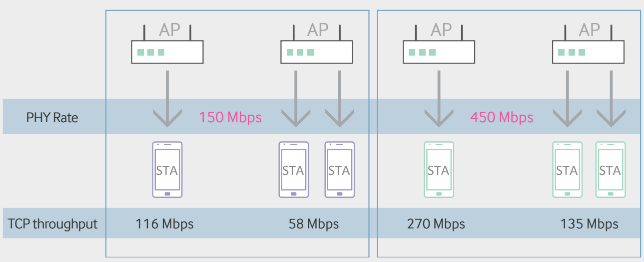

However, as it has been pointed out in the section above, the fairness principle has to be dealt with carefully, especially heterogeneous scenarios. The wireless resource distribution mechanisms have to make sure there is not a limitation and restricting the enhancements of faster clients and providing clients with the amount of airtime resources they are entitled to and preventing some stations from clogging the channel. Failure to achieve this fairness would lead to the starvation of some stations which would eventually degrade the overall performance of the network.
This is the reason why different approaches to the wireless resource distribution principle are studied:

- Throughput fairness: Providing throughput fairness effectively causes the clients to be entitled to the same effective amount of data transmission over a period of time. This is the type of fairness achieved with equal transmission opportunities in the 802.11 MAC when AMPDU is not present.

Mixed; Throughput fairness

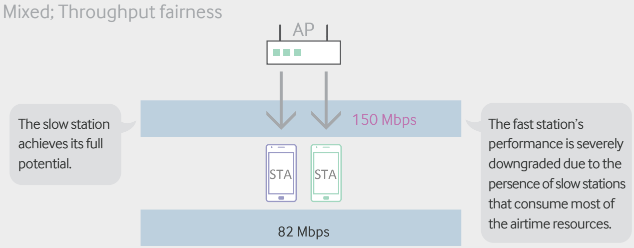

- Throughput maximization: The rationale behind this technique is quite simple; allow access to the airtime resources under the condition that the fastest transmissions are always allocated first. The presence of fairness in this technique is very limited since it may cause slow clients to starve if the number of fast clients is high, as they would always be granted access first.

Mixed; Throughput maximization

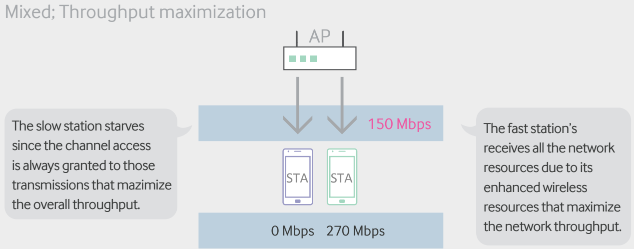

- Airtime Fairness: The concept is very simple but very powerful. Every client at a given quality-of–service level has the same right to use airtime. Therefore, the same amount of airtime resources should be allocated to each of them. Different clients can make use of their airtime according to their features, allowing enhanced clients with faster transmission rates to achieve higher throughput as their assigned airtime resources allow.

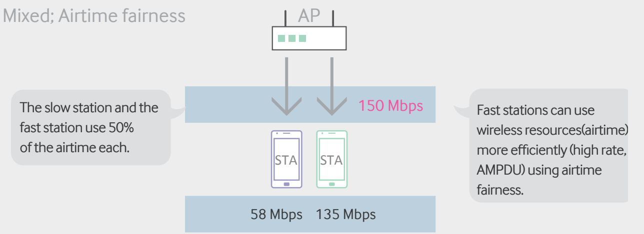

Clearly, a policy guaranteeing equal throughput or throughput maximization is not appropriate for wireless enterprise networks. Airtime fairness is essential to support critical enterprise applications since it prevents slow clients from slowing down the entire network or letting the fastest clients highjack the medium leading to the starvation of slower clients.
Despite clear benefits and the simplicity of the concept, airtime fairness is a not-so-simple algorithm and is rather difficult to put into practice. The implementation in enterprise wireless networks presents several challenges mainly due to the lack of knowledge by the network to allocate airtime fairly. This is the reason why Samsung Electronics designed the AirEqualizer algorithm to give its wireless enterprise APs an enhanced performance, guaranteeing the airtime fairness both for fast clients and legacy ones.

AirEqualizer

The most important goal of the AirEqualizer (AE) algorithm, which belongs to the Samsung Downlink Service, is to provide airtime fairness and avoid the monopolization of the AP service by a certain group of stations in the network that would result in unfair provision of wireless resources, starvation and unsatisfactory performance. In this section, the operation of AirEqualizer will be summarized to illustrate how it manages to achieve airtime fairness in an efficient way.

Enterprise WLANs have to deal with different types of traffic, from voice and video to best effort applications such as web browsing or emails. Each traffic type has very different requirements, and should therefore be treated differently. The channel access mechanism used in the enterprise network has to guarantee that all the traffic requirements are met so that the clients obtain a satisfactory level of quality of service (QoS). Best effort traffic, due to its nature and lower priority is less sensitive to network delays and jitter parameters or even loss rate. However, due to the greedy nature of best effort traffic and the fact that transmission schemes change dynamically according to location and medium conditions, best effort traffic is very prone to inefficiencies and performance degradation. It is for these particular reasons that the airtime fairness efforts of AirEqualizer focus mainly on best effort traffic. Unlike existing downlink AP scheduling approaches that use a single First-In First-Out queuing system for all the traffic for its clients, AirEqualizer provides differentiated services, per station, according to the priority of the data traffic using different scheduling algorithms:

- First-In First-Out queuing (FIFO): FIFO queues are the simplest and most common type of queue, packets in the queue are processed in the order they arrived. An illustrative example of FIFO queues would be the supermarket queues for clients that need to pay for the products they want to purchase. The clients that arrive first to the till are served first and therefore leave the queue in order of arrival.

AirEqalizer

- Weighted Round Robin Queuing: Each station queue is assigned equal time periods to process its packets and in a circular order. In order to give priority to some clients in the network weights can be included that increase the station processing period.

- Weighted Fair Queuing: Theoretical fair queuing is similar to round-robin bitwise scheduling amongst queues. The practical concept of the fair queue is to schedule packets destined to stations guaranteeing that all stations will effectively be served the same amount of resources taking into account the data packet sizes that need to be transmitted. Effectively, weighted fair queuing associates a weight to each packet of length (L). To make the scheduling selection proportional to the PHY rate this weight is chosen to be the transmission rate (R).
  In practice, the approached followed to implement is the start time fair queuing. The fair queuing selects the transmission order for the packets by keeping two tags the start time tag and the modeled finish time for each packet. Eventually, the packets are scheduled in the increasing order of the start tags of the packets, in other words the smallest starting round number is the next transmitted.
  An illustrative example of fair queuing is shown below, where we can see different queues competing for access channel and the WFQ scheduler that grants channel access calculating the amount of resources consumed by each of the STAs and allocating the same amount of airtime resources to each one of them.

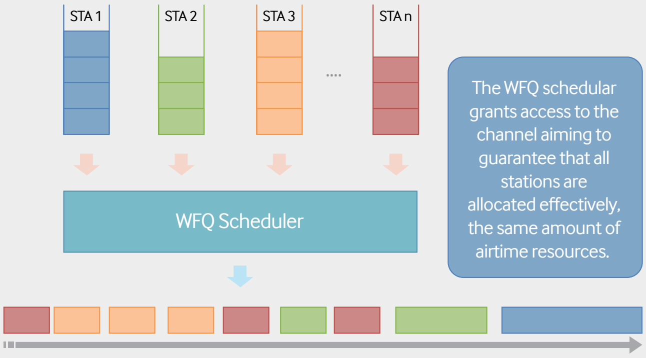

The AirEqualizer is shown in the Samsung Downlink Service structure diagram below, where the schedulers dealing with best effort and background traffic. We can observe that unlike other APs that use FIFO queues for all their traffic, Samsung AE utilizes WFQ:

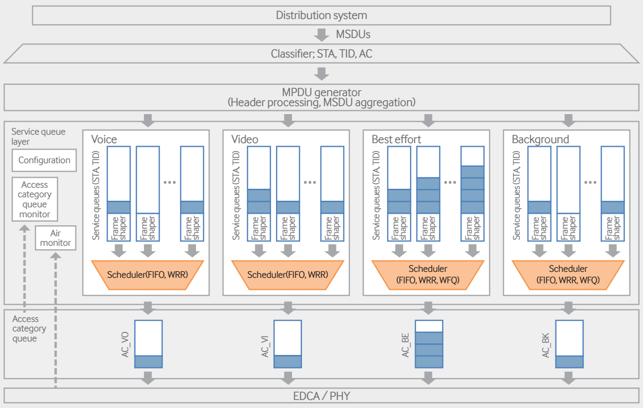

Air Equalizer structure
for downlink service

So the detailed steps of the AirEqualizer mechanism are the following:

- AirEqualizer classifies the packets that need to be transmitted from the AP to the stations according to the packet Access Category (AC) and Traffic Identifier and they are enqueued in the corresponding station queue they are destined to. Unlike the general structure, AirEqualizer has different station queues for each of the four AC types of traffic.

- AirEqualizer selects the station queue to serve based on the schedule appropriate schedule policy. As explained before, arrival time is not the only decisive factor but the amount of airtime resources used to serve the stations. Higher priority traffic (voice and video traffic), which have consistent, lower size frames are scheduled following the FIFO techniques or an optional weighted round robin. However, best effort traffic, whose data frames are larger and therefore require larger amounts of airtime resources, is scheduled following the weighted fairness queuing mechanism, aiming to allocate the same airtime resource to each receiving station.

- Eventually, the packet selected by the scheduler is enqueued in the downlink transmission AC queues of the AP and starts the CSMA/CA procedure to access the channel and finally be transmitted to its destined station.

AirEqalizer

- Throughout the transmission process, AirEqualizer monitors the consumption of airtime to serve the stations and uses the airtime as feedback to split the network resources evenly.

Ultimately, the use of the AirEqualizer algorithm achieves a fair and efficient downlink transmission that results in all the stations being served fairly according to their potential. Consequently, enhanced QoS for the network clients are achieved.

Results

In this section, a series of examples with the test results carried out using AirEqualizer is shown in order to illustrate the achievement of airtime fairness by the Samsung AirEqualizer algorithm and its direct impact on the improvement of the network performance.

Initially, the results of the Samsung AP with 802.11n stations are shown compared to equivalent AP models from other enterprise WLAN vendors. The performance metrics of interest in our case are the airtime fairness obtained by the AirEqualizer algorithm and the total TCP goodput achieved.

In order to illustrate the example let's think of a common enterprise where mobile stations share the airtime resources. In this particular example, clients are IEEE 802.11n stations located near the AP that receive best effort traffic from TCP applications with packets with a maximum segment size (MSS) of 1410 Bytes, using the 5 GHz band to transmit and receive packets. Some of the stations are slow 1-antenna smartphones and the others, fast 3-antenna notebook PCs capable of using 3x3 MIMO techniques.

Let's start our illustrative example with a case scenario in which thirty clients are present in the network. Initially, the vast majority of clients (i.e., 27 clients out of 30) are fast 3-antenna notebook PCs with good medium conditions and therefore able to use their fastest transmission schemes without compromising their performance (450 Mbps) whereas the rest of the stations are slow smartphones that use 150 Mbps to transmit their packets. The following two graphs show the experimental results obtained from such scenario using AirEqualizer in the Samsung AP and two other AP models from the competition.

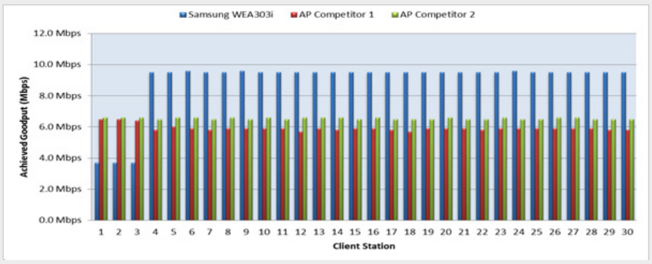

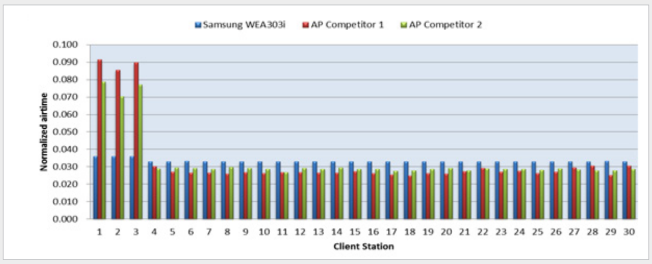

The first graph illustrates the TCP goodput achieved by each of the clients in the station. The competing APs present a uniform goodput performance, serving all of the stations an average of 6.5 Mbps regardless of their characteristic. This approach achieves a total goodput of 195 Mbps approximately. On the other hand, the Samsung AP differentiates between the clients with different characteristics. The slow smartphones are all served 3.5 Mbps while the faster notebook PCs achieve around 9.5 Mbps. An overall goodput performance of 267 Mbps is achieved in the WLAN using AirEqualizer. We can see that the first three clients, which represent the 1-antenna smartphones, cause a considerable degradation in the fast client performance when throughput fairness is used instead of the airtime fairness provided by AirEqualizer. Taking a look at the second graph, we can find the reason behind the first observation; those slow transmissions consume longer amount of airtime resources and prevent faster clients from profiting from their higher transmission rates.

AirEqalizer

Let's imagine that some of the notebook PC stations are replaced by more 1-antenna smartphones in the enterprise office. This effectively increases the number of slow stations and makes the scenario a 50% fast notebook PC clients and 50% slow smartphone clients. These are the results obtained in a situation as such:

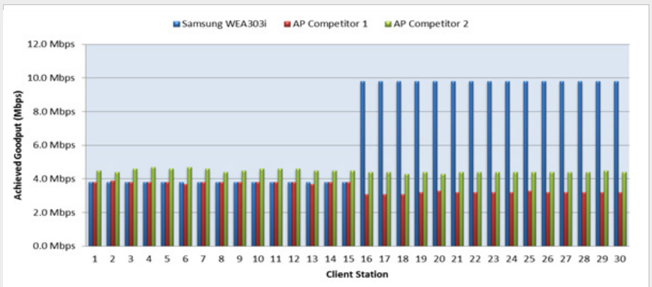

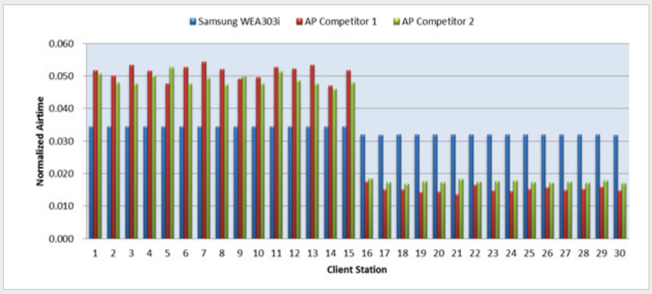

The increased presence of slower stations, damages the performance of those fast ones even more, and consequently the overall performance of the network is degraded. Unlike its competition, that achieves a total of 135 Mbps approximately, the Samsung AP is able to efficiently utilize the fast speed of the 3-antenna notebook PCs and achieve an overall throughput of 207 Mbps. AirEqualizer maintains a fair distribution of the airtime resources and therefore avoids the performance anomaly problem, making the performance of each client of the network independent of the transmission speed of its neighbours.

Let's imagine for a minute a drastic situation, where most of the clients are slow smartphones while three fast notebook PCs remain in the network. These are the results observed:

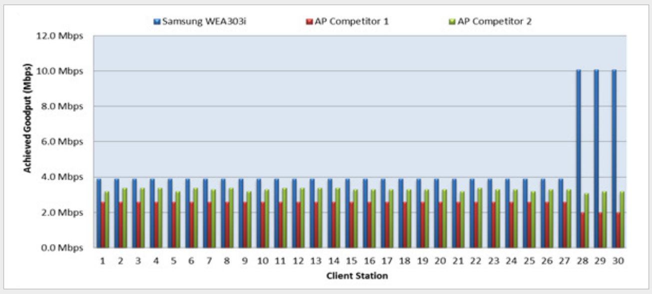

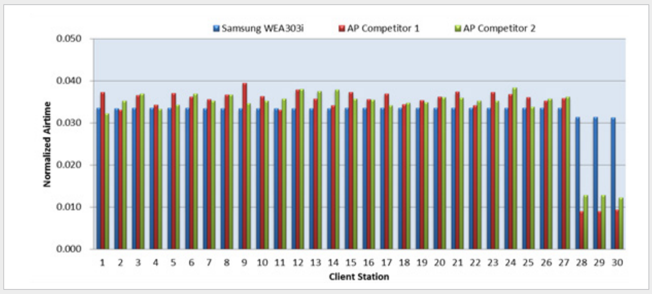

We have illustrated how AirEqualizer maintains the airtime fairness and consequently, the performance of each client is only dependent on its features and the total number of clients in the network. Changes in location or transmission scheme by any of the neighbouring clients does not affect another client's performance. We can consequently claim that the overall operation is fair. These experiments aim to demonstrate the AirEqualizer's ability to maintain the fairness of the airtime resources distribution despite the presence of fast and slow clients in the same environment. This is a summary of possible scenarios:

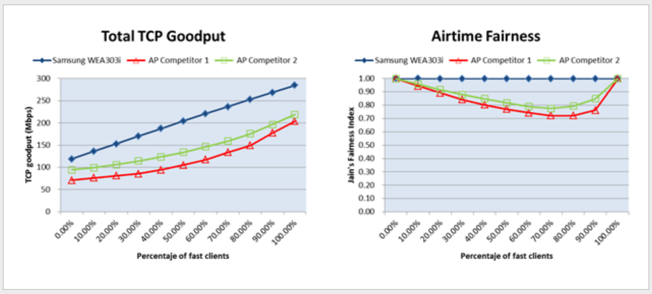

Samsung's AirEqualizer solution shows a noticeably higher TCP goodput consistently, regardless of the percentage of slow clients in the network. What's more, the airtime fairness index is always at its optimal point since the AirEqualizer service prevents any kind of unfair airtime appropriation.

Conclusion

After all the explanations and test results shown in this paper we can conclude that AirEqualizer is a more than adequate algorithm for enterprise WLANs. We summarize its benefits as follows:
Value for the clients

- AirEqualizer allows airtime fairness where all the stations in the network are allocated the same amount of airtime resources.
- No starvation, unfair resource appropriation or performance limitation due to inefficient network resource allocation.
- Stations are served the same amount of airtime resources used to the best of their capabilities.

Competitive strengths

- AirEqualizer achieves airtime fairness:
- Throughput increased over 60%
- Provides optimum wireless service by allocating equal airtime to each device.
- AirEqualizer maintains the benefits of having a multi-rate environ-ment.