



Providing QoS in WLANs

How the IEEE 802.11e Standard QoS Enhancements Will Affect the Performance of WLANs

This white paper discusses the key Quality of Service (QoS) and Medium Access Control (MAC) enhancements of the 802.11e standard that will affect performance in WLANs, and the industry readiness to adopt these key features.

Contents

Executive Summary	2	New Acknowledgment Rules	5
Introduction	2	Direct Link Protocol	5
Original 802.11 MAC	3	Piggybacking	5
IEEE 802.11e Standard	3	Automatic Power Save Delivery	6
Enhanced Distribution Coordinate Access (EDCA)	4	Optimizing for different Traffic Types	6
HCF Controlled Channel Access (HCCA)	4	Wi-Fi Alliance* Development of WMM and WMM-Scheduled Access Certifications	6
Traffic Specifications (TSPECs)	4	Industry Readiness for Adopting QoS Standards	6
802.11e MAC Enhancements	4	Conclusion	8
Contention Free Bursts	5		

Executive Summary

In wired Ethernet, Quality of Service (QoS) has been neglected as the bandwidth is very high and packet-error rates are very low. On the contrary, the wireless medium has relatively limited bandwidth and higher packet-error rates with high packet overheads. This could potentially limit the use of Wireless Local Area Networks (WLANs) for delivering traffic for real-time applications such as Voice over IP (VoIP) telephony and multimedia applications. End-users expect not only the mobility provided by WLANs but also QoS support. The upcoming IEEE 802.11e standard will add QoS to 802.11. This white paper presents some of the key QoS features and Medium Access Control (MAC) enhancements that are expected to be added by the 802.11e standard. Furthermore, it discusses the industry readiness to adopt some of the key components of the 802.11e standard.

Introduction

One of the drawbacks of wireless networks in comparison to wired networks is that they are generally less efficient and unpredictable. Wireless has limited bandwidth, high packet overheads, and is more prone to environmental factors such as obstructions, interference, weather and so on. The wireless medium (air) is much harder to control than a physical wire. The WLAN medium is also unlicensed and is therefore subject to interference from other devices. To further compound the

problem, wireless devices are generally constrained by size, weight and battery size, limiting the processing power and the battery life. These factors further limit the capability of the network to provide an optimal solution. The main objective of WLAN QoS is to optimize use of limited bandwidth offered by a WLAN to address the issues noted above.

To optimize the best use of the resources and fulfill the resource requirements of different applications, QoS provides mechanisms to control access and usage of the medium based on the application. Each application has different needs in terms of latency, bandwidth and packet-error rate and, therefore, QoS must cater to each of these needs. Applications requiring low latency (e.g., voice) may be given higher priority to use the medium, whereas applications requiring higher bandwidth may be assigned longer transmit times (e.g., video). Other traffic may require high reliability (e.g., email and data) and must be delivered with low packet-error rate.

The original 802.11¹ standard was not designed to provide differentiation and prioritization based on the traffic type, thus providing less than optimal user experience for voice and video over WLAN applications. Voice applications require no dropped calls or bad connections. Video/audio applications

¹ IEEE Std 802.11-1997 - Information Technology- Telecommunications And Information Exchange Between Systems-Local And Metropolitan Area Networks-specific Requirements-part 11: Wireless LAN Medium Access Control (MAC) And Physical Layer (PHY) Specifications.

require enough bandwidth to maintain high quality video/audio streams. Email and file-sharing applications require ensuring delivery of error-free files. To fulfill these requirements, the upcoming IEEE 802.11e standard will add several QoS features and enhancements to WLAN. The key benefits of the 802.11e standard are:

- Reduces latency by prioritizing wireless packets based on traffic type.
- Enables Access Point (AP) to schedule resources based on client/station data rate and latency needs.
- Improves wireless bandwidth efficiency and packet overheads.

The 802.11e standard provides a comprehensive QoS solution and will contain many optional features. The industry may not yet be ready to adopt all the optional features. In this white paper we discuss key features of the 802.11e standard, their benefits and industry readiness to adopt these features.

Original 802.11 MAC

The original 802.11 MAC does not provide differentiated services based on traffic type. However, as wireless networks provide multimedia services involving voice and video, high packet overhead with limited bandwidth in a WLAN can become a major stumbling block for delivering delay-sensitive packets. In the original 802.11 standard, as much as one third of the data rate can be consumed by packet fragmentation, inter-frame spacing and acknowledgments. Furthermore, under heavy traffic load conditions, collisions and backoffs can severely deteriorate the quality of voice and video applications.

The 802.11 standard specifies two channel access mechanisms: Distributed Coordination Function (DCF) and Point Coordination Function (PCF). DCF allows sharing of the wireless medium between the Stations (STAs) and the AP using Carrier Sense Multiple Access with Collision Avoidance (CSMA/CA). DCF provides best-effort service and does not provide either medium access priority or support delay and bandwidth requirements of different applications. DCF mode operates as follows. Each station checks whether the medium is idle before transmitting. If the medium is detected to be idle for Distributed Inter-frame Space (DIFS) interval of time, the station begins transmission. In the case where the medium is determined to be busy, the station defers transmission until the medium is idle for DIFS time. The station then selects a random backoff interval (using a backoff algorithm) and decrements a backoff counter while the medium is idle. The backoff mechanism is used to prevent two or more stations from

transmitting simultaneously. Once the backoff interval has expired, the station begins its transmission. The range from which the random backoff interval is selected is called the Contention Window (CW) and depends on the number of previous retransmission attempts. Once the MAC Service Data Unit (MSDU) has been transmitted, the station waits for a Short Inter-Frame Space (SIFS) duration for the Acknowledgment (ACK) from the recipient.

PCF is an optional channel-access mechanism in the 802.11 standard that is not commonly implemented due to lack of market demand. PCF provides contention-free access to the medium. It was designed to support time-sensitive applications. The Point Coordinator (PC) residing in the AP provides a contention-free wireless medium access. A polling method is used to provide access, with the PC acting as the polling master. This eliminates collisions and the time spent on backoff and contention as described previously for DCF.

Contention-free access to the medium is not provided at all times. When PCF is used, time on the medium is divided into a Contention-Free Period (CFP) and a Contention Period (CP) by the PC. During CFP and CP, PCF and DCF are used to access the medium, respectively.

Neither DCF nor PCF has sufficient functionality to provide QoS demanded by multimedia applications. DCF treats all traffic the same—with all stations contending for the medium with the same priority. PCF also has several inadequacies in the support of QoS:

- Lack of mechanisms to differentiate different traffic types.
- No mechanisms for the stations to communicate their QoS requirements to the AP.
- No management interface to control and setup CFP.
- The polling schedule is not tightly controlled.

In the next section, we discuss how the 802.11e standard addresses these drawbacks of 802.11 to provide QoS.

IEEE 802.11e Standard

The IEEE 802.11 Task Group E (802.11e) has defined enhancements to the original 802.11 MAC (Medium Access Control) to provide QoS. The 802.11e standard introduces the Hybrid Coordination Function (HCF), which combines functions from DCF and PCF with enhanced QoS-specific mechanisms and frame types. HCF has two modes of operation—Enhanced Distribution Coordinate Access (EDCA) and HCF Controlled Channel Access (HCCA). EDCA and HCCA are contention-

based and polling-based mechanisms for channel access, respectively, and operate concurrently. During the CP, EDCA is used for channel access; whereas during CFP, HCCA is mostly used. A Station (STA) that supports QoS is referred to as a QoS Enhanced Station (QSTA); whereas an Access Point (AP) that supports QoS is referred to as a QoS Enhanced AP (QAP).

HCF allocates QSTAs the right to transmit through Transmission Opportunity (TXOP). A TXOP defines the start time and the maximum duration during which a QSTA can transmit a series of frames. A summary of EDCA and HCCA mechanisms are provided in the sections of this paper titled “Enhanced Distribution Coordinate Access (EDCA)” and “HCF Controlled Channel Access (HCCA),” respectively.

Enhanced Distribution Coordinate Access (EDCA)

EDCA contention access is an extension to DCF and provides prioritized access to the wireless medium. The EDCA channel-access mechanism defines four Access Categories (AC) based on the IEEE 802.11 standard² to provide priorities. Each AC has its own transmit queue. The following four key parameters are used for differentiation:

- Minimum contention window size (CW_{Min}). AC with higher priority is assigned a shorter CW_{Min} .
- Maximum contention window size (CW_{Max}).
- TXOP limit—Specifies the maximum duration a QSTA can transmit and is specified per AC. The TXOP limit can be used to ensure that high-bandwidth traffic gets greater access to the medium. TXOP limit also makes the channel-access protocol significantly more efficient.
- Arbitration Inter-Frame Space (AIFS)—specifies the time interval between the wireless medium becoming idle and the start of channel-access negotiation. Each AC is assigned a different AIFS[AC] based on the AC to further provide QoS differentiation.

Each AC contends independently for TXOPs based on the above parameters within the QSTA. Once the AC has sensed that the medium has been idle for AIFS[AC], it starts its backoff time (similar to DCF). If there is a collision between ACs within a QSTA, data frames from the AC with the highest priority receives a TXOP. Data frames from the remaining ACs behave as if there was an external collision.

² IEEE Std 802.11-2004 (Revision of IEEE Std 802.11-1998)—IEEE Standard for Local and metropolitan area networks Media Access Control (MAC) Bridges.

HCF Controlled Channel Access (HCCA)

HCCA uses a Hybrid Coordinator (HC) to centrally manage the wireless medium access to provide parameterized QoS. Parameterized QoS refers to the capability of providing QoS flows from applications with specific QoS parameters—such as data rate, latency and so forth. Like PCF, HCCA uses a polled-based mechanism to access the medium, thereby reducing contention on the wireless medium. The key differences between HCCA and PCF are that HCCA can poll the stations during CP and that it supports scheduling of packets based on QSTA's specific traffic-flow requirements. The traffic-flow requirements of the QSTAs are specified using Traffic Specifications (TSPECs) discussed in the section of this paper titled “Traffic Specifications (TSPECs).”

The HC has the highest priority over all QSTAs in gaining access to the medium as it has the shortest waiting time compared to all QSTAs backoff times. The HC provides contention-free frame exchange with short delays, thereby providing tighter controlled latency.

Traffic Specifications (TSPECs)

Because WLANs have limited bandwidth and provide contention-based medium access, they are susceptible to traffic congestion, which can lead to severe overall network performance degradation. As the networks become overloaded, CW sizes can increase significantly, leading to long backoff times. This can become a network performance bottleneck as the limited bandwidth of the network is not fully utilized. This necessitates some admission-control mechanism to be built into the standard for regulating traffic. The IEEE 802.11e standard specifies the use of TSPECs for negotiating admission control for both EDCA and HCCA.

TSPECs are used by QSTAs to specify traffic-flow requirements such as data rate, delay, packet size and service interval. The QAP may accept or reject a new TSPEC request based on the network conditions. If a TSPEC is rejected by the QAP, high-priority AC inside the requesting QSTA is not allowed to use high-priority access parameters.

802.11e MAC Enhancements

Apart from providing EDCA and HCCA functionality discussed previously, the 802.11e standard provides several MAC enhancements. Some of the key enhancements are summarized next.

Contention Free Bursts

Contention Free Bursts (CFB) allow a QSTA/QAP to send several frames in a row without having to contend for the medium again and again. The QSTA/QAP continues to transmit after the SIFS delay if there is time remaining in a granted TXOP. CFBs may be used during TXOPs that were gained during EDCA or HCCA. Figure 1 shows data transmission with and without CFB. Bursting can significantly improve performance as the overheads associated with DIFS and backoff are reduced.

CFB can also be used to improve 802.11g's throughput in a mixed 802.11b and 802.11g environment as it frees time for 802.11g STAs. A fast 802.11g STA can transmit several frames in the time period that a slower 802.11b STA takes to transmit one frame and its associated overhead. An 802.11g STA can be assigned a TXOP comparable to the single-frame duration of a STA that uses 802.11b.

Some vendors have implemented proprietary packet-bursting schemes. The 802.11e standard enables a standardized method of implementing bursting, providing interoperability across multiple vendors' network gear.

Figure 1.a Data Transmission with No Bursting

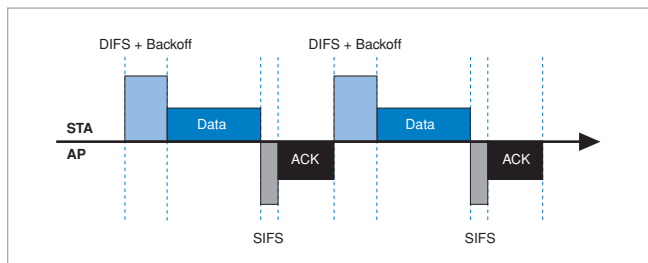
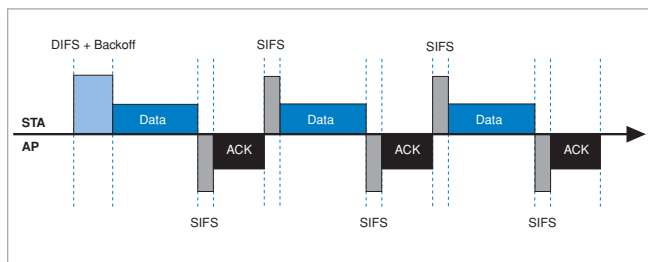


Figure 1.b Data Transmission with Bursting



New Acknowledgment Rules

In the 802.11 standard, all unicast data frames require an immediate ACK control frame. HCF adds these two new options that are specified in the QoS control field of data frames:

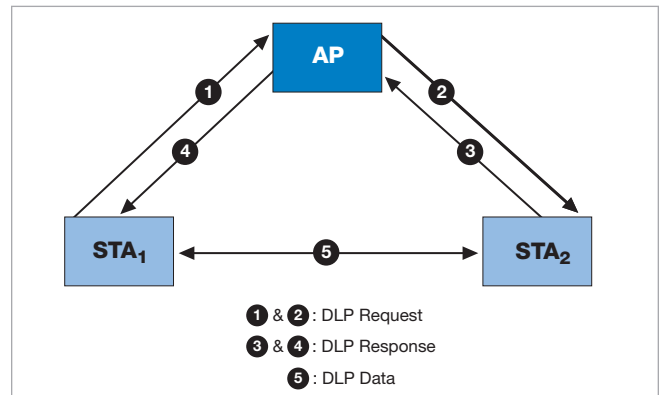
- No ACK—Increases efficiency by sending no ACK for certain applications. This feature is useful for applications with very low latency tolerance, but which can tolerate a significant amount of packet loss (e.g., Voice over IP).

- Block ACK—Increases efficiency by aggregating ACKs for multiple frames into a single response. Block acknowledgements are of two types:
 - Immediate—In which the receiver must respond immediately with a block ACK frame once it receives a block ACK request frame from the sender. The block ACK request is sent by the sender after transmitting multiple data frames in a CFB during the TXOP.
 - Delayed—In this case the receiver responds to the block ACK request frame from the sender indicating that the block ACK will be delayed. This provides more time for calculating the ACK for lower-performance systems.

Direct Link Protocol

The original 802.11 specification permits traffic in an AP-based network to flow between the STAs and AP only. Direct Link Protocol (DLP) in the 802.11e standard adds the capability for the STAs to send traffic directly to each other without traversing the AP, as shown in Figure 2. This capability can potentially increase available bandwidth for STA-to-STA communication by possibly two times. DLP will function only when the communicating STAs are in range of each other. DLP could also potentially increase the performance in the case where the link performance between the communicating STAs is better than the link performance between the STAs and AP. This could be the case when the STAs are closer to each other than to the AP.

Figure 2 Data Link Protocol



Piggybacking

Piggybacking reduces polling and ACK overhead by sending data “piggybacked” on polls and on ACKs. This feature can improve the overall network performance.

Table 1: Applications Requirements

Application	Bandwidth	Latency	Frame Loss
Audio	64 Kbs - 1.5 Mbps	Higher tolerance to latency	Higher tolerance to frame loss
Video streaming	Higher requirement: 2 Mbps for SD and 20 Mbps for HD	Low	Higher tolerance to frame loss
Voice and Video conferencing	Lower requirement: voice < 32 Kbps; 128 Kbps for videoconferencing	Low (<50 ms)	Higher tolerance to frame loss
Gaming	Lower requirement: 32 Kbps – 128 Kbps	Lowest (< 10 ms)	Low tolerance to frame loss

Automatic Power-Save Delivery

Automatic Power-Save Delivery (APSD) enables extended battery life by allowing devices to turn off their radios for most of the time. It is an enhancement to the existing 802.11 power-save mechanism. APSD enables a station to set up a “schedule” for delivery of frames, based on a repeating pattern of a specified number of beacon intervals. This feature would be appropriate for VoIP applications where packets are transmitted on a regular interval.

Optimizing for Different Traffic Types

Table 1 shows different application requirements in terms of bandwidth, latency and frame-loss tolerance. Applications requiring higher bandwidth require longer TXOP so that they have an opportunity to transmit more data. Bursty traffic such as File Transfer Protocol (FTP) have a high-bandwidth requirement but do not necessarily require low latency. Thus, they need to be assigned longer TXOP. Latency- and jitter-sensitive applications require consistent timely access to the medium; therefore, they should be assigned higher priority to access the medium. Voice requires low latency but not high bandwidth. Data traffic may not need low latency or extremely high bandwidth, but traffic must get there, and therefore requires high reliability.

In applications with very low jitter tolerance (e.g., streaming media), where data may no longer be useful after a retry, ACK is not important. On the other hand, in applications requiring high reliability (e.g., data traffic), ACK is important to improve reliability.

Wi-Fi Alliance* Development of WMM and WMM-Scheduled Access Certifications

Wi-Fi Multimedia (WMM) is a subset of the draft 802.11e standard and is motivated by the need to provide an early QoS implementation. WMM is also known as Wireless Multimedia Enhancements (WME). The Wi-Fi Alliance (WFA)*

developed requirements and carries out interoperability testing to ensure products from different vendors are interoperable. The key focus of WMM is the EDCA portion of the 802.11e standard. WFA certification for WMM interoperability is expected to begin in September 2004.

WMM-Scheduled Access, also known as Wi-Fi Scheduled Multimedia (WSM), is a subset of the draft 802.11e standard. The key feature of WMM-Scheduled Access is HCCA. WMM-Scheduled Access certification is expected to begin in the first quarter of 2005.

Industry Readiness for Adopting QoS Standards

Though the 802.11e standard is expected to be ratified by the end of 2004, several of the components of the 802.11e standard might not be adopted by the industry in the near future. This is due to the lack of QoS infrastructure and interoperability testing.

The QoS infrastructure not only includes STAs and APs that support QoS but also operating systems (OS) and applications supporting QoS. Current OS and applications provide support for WMM/EDCA but lack support for WMM-Scheduled Access/HCCA. This could be one of the hurdles for an early industry-wide adoption of WMM-Scheduled Access.

Furthermore, different implementations of WMM and WMM-Scheduled Access among vendors could cause lack of interoperability. Wi-Fi Alliance certification for WMM interoperability is expected to begin in September 2004; thereafter, interoperability may not be an issue. On the other hand, WMM-Scheduled Access certification is not expected to begin before the first quarter of 2005, delaying its wide adoption.

One of the other issues that could get in the way of early WMM-Scheduled Access adoption is fast roaming. As a QSTA roams from one QAP to another, the old and new QAPs should be able to coordinate their traffic-flow schedules

without requiring the QSTA to renegotiate its traffic-flow requirements with the new QAP. Traffic-flow schedule coordination between QAPs is not addressed by the 802.11e standard and would be required for WMM-Scheduled Access to enable fast roaming. This would require QAP vendors to develop proprietary solutions beyond those of the 802.11e standard. This could potentially raise interoperability issues for providing fast QoS roaming, further delaying wide adoption of WMM-Scheduled Access.

This raises the question whether WMM is “good enough” to provide QoS and removes the motivation to deploy the required WMM-Scheduled Access infrastructure. Are there benefits of WMM-Scheduled Access which can offset the cost of the infrastructure? The answer, of course, comes down mainly to performance and complexity, the wireless environment in which the wireless network resides, and the applications.

Provisioning of QoS in a wireless environment can be a real challenge for several reasons:

- Interference from other wireless sources could cause severe degradation in the wireless link performance, jeopardizing QoS.
- The WLAN may be supporting QoS aware (QSTA) and non-QoS aware (STA) stations, complicating QoS provisioning.
- All the traffic may not be well behaved (i.e., fits normal probability distribution).

Thus, neither WMM nor WMM-Scheduled Access may be able to guarantee QoS. We now compare WMM and WMM-Scheduled Access to answer the above question.

Because WMM-Scheduled Access uses a central reservation scheme, the complexity of the reservation scheme may negatively impact network performance. However, the network performance degradation may be offset by the QoS benefits for certain applications, though not all applications may realize the QoS benefits. WMM may be sufficient to provide QoS in environments where contentions occur infrequently and, therefore, little benefit would be harnessed from the centrally managed wireless medium access provided by WMM-Scheduled Access. When there are other WLANs within range (e.g., in the neighbor’s house) contending for the same bandwidth, WMM may provide performance benefit over WMM-Scheduled Access. From the applications point of view, WMM-Scheduled Access may be more suitable for applications requiring constant bit rate (e.g., VoIP) whereas WMM may be more suitable for applications requiring variable bit rate (e.g., video). Hence WMM-Scheduled Access may be more widely deployed on handsets in the near future.

Table 2 below summarizes the key differences between WMM and WMM-Scheduled Access. Until a proper infrastructure for WMM-Scheduled Access is in place, WMM can be expected to play a key role in providing QoS for a WLAN.

Table 2: WMM versus WMM-Scheduled Access

Parameters		WMM	WMM-Scheduled Access
Channel Access Mode		<ul style="list-style-type: none"> • Contention based • Distributed control 	<ul style="list-style-type: none"> • Poll based • Central reservation scheme
Key Technical Benefits		QSTAs and QAPs can provide prioritization among wireless data streams to distinguish between high and low priority traffic	Provides QAP the mechanism to centrally reserve bandwidth: <ul style="list-style-type: none"> • QAP can refuse a connection when not sufficient bandwidth is available • QAP can ensure that subscribed % of media access will be given to the subscribers.
Applications Dependency		Current applications support WMM as WMM uses 802.1D priority and the applications can be agnostic	Current applications do not support WMM-Scheduled Access as the applications have to be wireless aware
OS Dependency		Current OS support provided for WMM	No current support for WMM -Scheduled Access
Fast Roaming Support Requirements		No requirements	Traffic flow schedule coordination between QAPs required and it is not supported in 802.11e
Usefulness	Environment	More appropriate to provide QoS in environments where: <ul style="list-style-type: none"> • Other WLANs are within range (e.g. in the neighbor’s house) 	More appropriate to provide QoS in environments where: <ul style="list-style-type: none"> • There are lots of devices • There are no other WLANs within range contending for the same bandwidth.
	Applications	More suitable for applications that require variable bit rate (e.g. video)	More suitable for applications that require constant bit rate (e.g. VoIP)
User intervention		No user intervention required	User intervention will be required for configuration—no standardized API
Interoperability Testing (WFA Certification)		Expected to begin in September, 2004	Expected to begin in first quarter of 2005

Conclusion

This white paper discussed some of the key QoS features and MAC enhancements in the upcoming 802.11e standard.

Though the 802.11e standard is expected to be ratified by end of 2004, not all the components of 802.11e (such as WMM-Scheduled Access/HCCA) will be widely adopted in the industry for an extended period of time—due to the lack of infrastructure availability. Some of the hurdles foreseen for the early adoption of WMM-Scheduled Access are infrastructure availability, user-intervention requirements and interference from other WLANs within the range.

As the infrastructure is already in place for WMM and as WMM is more applicable for applications today, it is likely that WMM will be the key component of the 802.11e standard for an early industry-wide adoption in the near future.



Intel and the Intel logo are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries.

*Other names and brands may be claimed as the property of others. This document is provided as a convenience and Intel does not warrant the accuracy, reliability or completeness of the information incorporated herein.

Copyright © 2004, Intel Corporation. All rights reserved. 0804/JM/MESH Please Recycle

303762-001US