Using IEEE 802.11e MAC for QoS over Wireless

Priyank Garg, Rushabh Doshi, Russell Greene, Mary Baker, Majid Malek, Xiaoyan Cheng
{priyank, radoshi, rdg12, mgbaker, mmalek, cheng042}@stanford.edu
Computer Science Department, Stanford University

Abstract

We study the behavior of the new MAC protocols for QoS in the proposed IEEE 802.11e draft standard and analyze them for their ability to fulfill their goals of better QoS and higher channel efficiency. We study the response of these mechanisms to various choices in available protocol parameters. We show that HCF reduces channel contention and allows better channel utilization. However, both the proposed MAC coordination functions, EDCF and HCF, are highly sensitive to protocol parameters. We believe that the effectiveness of these functions also depends on the scheduling algorithms. The effects of the various policy choices need to be understood and validated before the draft becomes a standard.

1 Introduction

The past few years have seen an explosion in the deployment of Wireless LANs (WLANs) conforming to the IEEE 802.11 standard. As a result, they are expected to support the same applications as the wired Ethernet that they are replacing. While QoS issues in Ethernet have been considered uninteresting due to huge improvements in the physical layer bandwidths, the IEEE 802.11e group is developing MAC improvements to support QoS sensitive applications, to enable a better mobile user experience and to make more efficient use of the wireless channel.

The original 802.11 MAC protocol included two modes of operation characterized by coordination functions: the Distributed Coordination Function (DCF) and the Point Coordination Function (PCF). The 802.11e draft specification introduces two new modes of operation, EDCF (Enhanced DCF) and HCF (Hybrid Coordination Function) which define mechanisms to enable QoS. What is not well understood in the specification are the effects of the various parameters and scheduling choices in each mode that give dramatically different results in various scenarios.

In this paper we compare the performance of the existing 802.11 MAC [?] and the proposed 802.11e draft standard [?] by simulating the protocols in NS2 [?]. Both EDCF and HCF are highly sensitive to the tuning parameters for the protocol. We study the response of the system to these parameters and compare the QoS provided by each mode of the new MAC. We find that they enable differentiated treatment of traffic streams and can be tuned to meet the QoS requirements of low latency and jitter. We also see that HCF enables more efficient use of the medium giving greater net throughput.

Sections 2 and 3 present an overview of the current 802.11 MAC and the QoS enhancements proposed. Section 4 discusses the various changes made to NS2.1b6 to simulate EDCF and HCF. In Section 5 we present the simulation environment used to obtain the results that are subsequently presented in Section 6. Section 7 describes future work and finally, we present our conclusions in Section 8.

2 Legacy 802.11 MAC

DCF is based on CSMA/CA (Carrier Sense Multiple Access with Collision Avoidance) In CSMA/CA, once a station detects that the medium is free it begins to decrement its back-off counter. Each station maintains a Contention Window (CW) that is used to determine the number of slot times a station has to wait before transmission. The back-off counter only begins to decrease after the medium has been free for a DIFS period (DCF Inter-Frame Space). If the back-off counter expires and the medium is still free the station begins to transmit. In case of a collision the station randomly picks a new back-off period from its CW, which grows in a binary exponential fashion like Ethernet, and then attempts to gain control of the medium again. Due to collisions and the binary back-off mechanism, there are no transmit guarantees with DCF.

With PCF, the period after each beacon transmission is divided into two sections, the Contention Free Period (CFP) followed by the Contention Period (CP), which together constitute a superframe. The point coordinator (PC, generally assumed to be co-located at the Access Point, AP) is guaranteed access to the medium in the beginning of the CFP. During the CFP, the PC lets stations have priority access to the medium by polling them in a round-robin fashion. Any station that requests to be added to the polling sequence is included during the next polling interval.

3 802.11e MAC Enhancements

The current MAC has no means of differentiating traffic streams or sources. All data is treated equally in both DCF and PCF. As a result, no consideration can be made for the service requirements of the traffic on the channel. For example, low priority bursty traffic may choke out a long-running critical video feed thereby destroying the user's experience.

The two new MAC modes, EDCF and HCF, being defined under 802.11e, support up to eight priority Traffic Classes (TC) that map directly to the RSVP protocol and other protocol priority levels [?] [?].

3.1 Enhanced Distributed Coordination Function (EDCF)

EDCF is DCF with some of the elements of the MAC parameterized per-TC. Each TC starts a back-off after detecting the channel being idle for an Arbitration Inter-Frame Space (AIFS), which is also referred to in this document as *CWOffset*. The AIFS is at least as large as the DIFS and can be chosen individually for each TC. It provides a deterministic priority mechanism between the TCs.

The minimum initial value of the CW, denoted by CWMin can be selected on a per-TC basis. As collisions occur, the subsequent CW is doubled (binary exponential back-off), thus providing a probabilistic priority mechanism between the TCs. The CWMax value sets the maximum possible value for the CW and is intended to be the same for all TCs as in DCF.

Within a station, the eight TCs have independent transmission queues. These behave as virtual stations with the above-mentioned parameters determining their ability to transmit. If the back-off counter of two or more parallel TCs in a single station reach zero at the same time, a scheduler inside the station treats the event as a virtual collision without recording a retransmission. The Transmit Opportunity (TXOP) is given to the TC with the highest priority of the colliding TCs, and the others back-off as if a collision on the medium occurred. These QoS parameters can be adapted by the Base Station via the beacon frames.

3.2 Hybrid Coordination Function (HCF)

HCF is an extension of the polling idea in PCF. As in PCF, under HCF, the superframe is divided into the CFP that starts with every beacon, and the CP. During the CP, access is governed by EDCF, though the Hybrid Coordinator (HC, generally collocated at the Access Point, AP) can initiate HCF access at any time.

During the CFP, the HC issues a QoS CF-Poll to a particular station to give it a TXOP, specifing the start time and maximum duration. No stations attempt to gain access to the medium at this time, and so when they receive a CF-Poll, they assume a TXOP and transmit any data they have. The CFP ends after the time announced by the beacon frame or by a CF-End Frame.

If a station is given a CF-Poll, it is expected to start responding with data within a SIFS (Short Inter-Frame Space) period. If it does not, the HC can take over the medium after a PIFS (PCF Inter-Frame Space) time, and allocate another CF-Poll to another station. This allows very efficient use of the medium during the CFP.

HCF also allows the base station to initiate a CF-Poll based transmission sequence within the CP, if it so desires.

3.3 Scheduling

The HC has available over time a snapshot view of the per-TC per-Station Queue length information in the cell, including that of the AP itself. This information is sent to the HC by the stations via the new QoS control field added to the MAC frame definition. With this, the HC has to decide to which station (including itself) to allocate TXOPs during CFP. At minimum, the following need to be considered: 1) Priority of the TC, 2) Required QoS for the TC (low jitter, high bandwidth, low latency, etc.), 3) Queue lengths per TC, 4) Queue lengths per station, 5) Duration of TXOP available and to be allocated, and 6) Past QoS seen by the TC.

When a wireless station receives a TXOP from the HC, the HC does not specify a particular TC for the TXOP leaving this decision up to the wireless station. This choice depends on the same factors as the HC Scheduler, except for the multi-station cell-wise aggregation that the HC scheduler uses. Decentralizing this decision allows a scalable mechanism of maintaining TC history and servicing TCs as per QoS seen in the past.

4 Implementation

We used Berkeley NS v2.1b6 to implement HCF. We updated Atheros Communications' EDCF code [?] on the same version of NS2 to conform to the latest draft specification. We added HCF frame exchange sequences. We also updated beaconing and added various new packet types.

We enhanced the stations (QSTAs) to transmit information regarding the states of their various queues to the AP. The AP in turn had to be able to understand this information and use it to provide TXOPs to the stations.

4.1 Schedulers

For the HC scheduler we currently implement a scheme that uses weighted average queue length per station. The weights are based on TC queues within a station. The scheduler allocates the maximum available TXOP within the CFP to the station with the largest average. Higher priority levels were given higher weights.

The End Point scheduler has access to the internal queues of the station as well as to TXOP information. Our scheme uses a first fit algorithm to find the highest priority packet that can be transmitted in the given TXOP.

5 Simulation

All scenarios were implemented in NS v2.1b6 for 20 seconds with a medium bandwidth of 11Mbps.

5.1 Scenario 1

The first scenario defines three nodes, an access point (AP) and two nodes contending for the medium. There are two traffic flows, a high priority 6 Mbps video stream and a background FTP stream with 1500 byte packets. The video stream also requires low jitter to minimize buffering requirements. We use this scenario to show how varying the different parameters in EDCF affect the traffic behavior.

5.2 Scenario 2

The second scenario has three stations and an access point. Station 1 is serving a 6 Mbps streaming video stream and a 34 Kbps audio stream to the access point. Stations 2 and 3 are also contending for the medium attempting to utilize 1.5 Mbps per flow for three flows providing some background traffic. The highest priority traffic class is the audio stream because it has the most stringent latency requirements as it represents an interactive conversation. It needs low round-trip latency, with a goal of less than 10ms. The video stream is a lower priority traffic class than the audio stream but higher than the background flows.

This scenario is more realistic than Scenario 1 as it has more participants. Here we have multiple contending entities within a cell, which performs poorly under DCF. Also, some of the stations have more than one type of originating traffic with different priority and QoS requirements.

6 Results

From our simulations we find that with the two parameters defined for EDCF, CWOffset and CWMin, traffic streams with more stringent latency and jitter requirements can be prioritized. However, the CWOffset is more influential than CWMin in achieving the desired results. We also observe that both EDCF and HCF enable meeting the QoS requirements for the scenarios we have defined. However, while EDCF does this without increasing overall efficiency of the medium, HCF gets up to 20% more throughput from the channel while still meeting the QoS constraints.

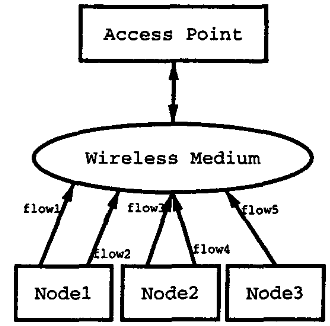


```
flow1:     34kbps  Audio Stream
flow2:     6Mbps   Video Stream
flow3,4,5: 1.5Mbps Average Rate
           Background Trafic
```

Figure 1: Simulation Scenario 2

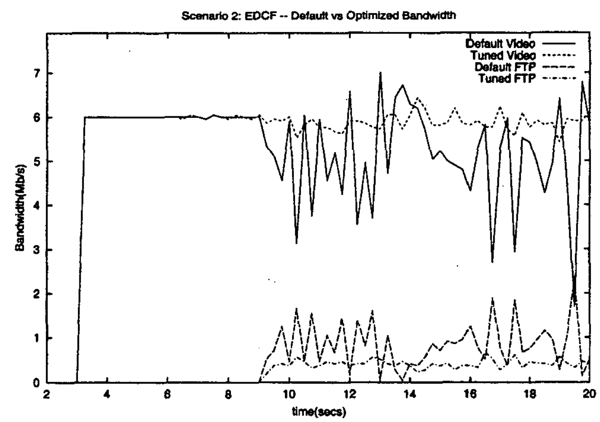

Figure 2: Tuned vs Default EDCF

6.1 Effect of EDCF Parameters

EDCF extends DCF by simply making some of the parameters of the CSMA/CA MAC protocol variable on a per-TC basis. However, the introduction of TCs which perform virtual collisions within a wireless endpoint and back-off accordingly affects the behavior of EDCF compared to DCF. Figure 2 compares the bandwidth performance of EDCF using the default values for CWOffset and CWMin for each TC with the performance after these parameters are tuned for best results. As seen in Figure 2, tuning EDCF significantly reduces the variation in bandwidth throughput.

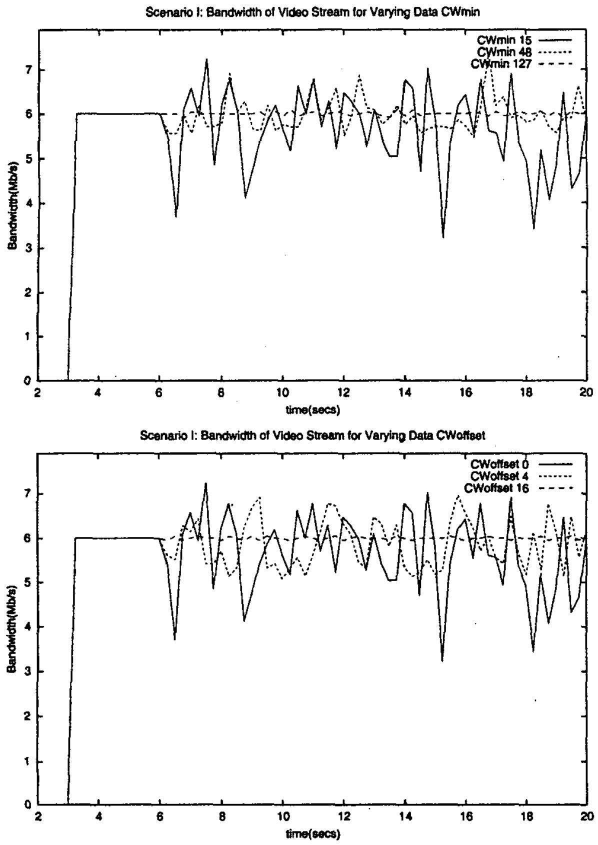

Figure 3: Tuning EDCF Parameters

Figure 3 and 4 show the bandwidth performance and the Cumulative Distribution Function (CDF) of latency for the video stream as the AIFS (CWOffset) and Minimum Contention Window (CWmin) are varied for the background traffic, the FTP stream. Increasing the CWOffset and CWmin of the FTP stream greatly reduces the bandwidth variance of the video stream and also significantly improves its latency and jitter characteristics. The highest latency for the video stream goes from 180ms to 40ms as the CWOffset is increased from a default of 0 to 16 for the FTP stream, or as CWMin is increased from a default of 15 to 127. At those values of CWOffset and CWMin for the FTP stream, the bandwidth for the video stream is near constant. Also, CWOffset has a greater effect on reducing bandwidth variance and jitter. The CDF for latency becomes steeper as CWOffset increases indicating that more packets have a latency within a small range of values, meaning jitter is reduced. CWMin, though reducing overall latency, does not increase the slope of the CDF, which means jitter is not improved. Reduced jitter allows smaller client buffers, thus cutting the wait time before a streaming video can start playing.

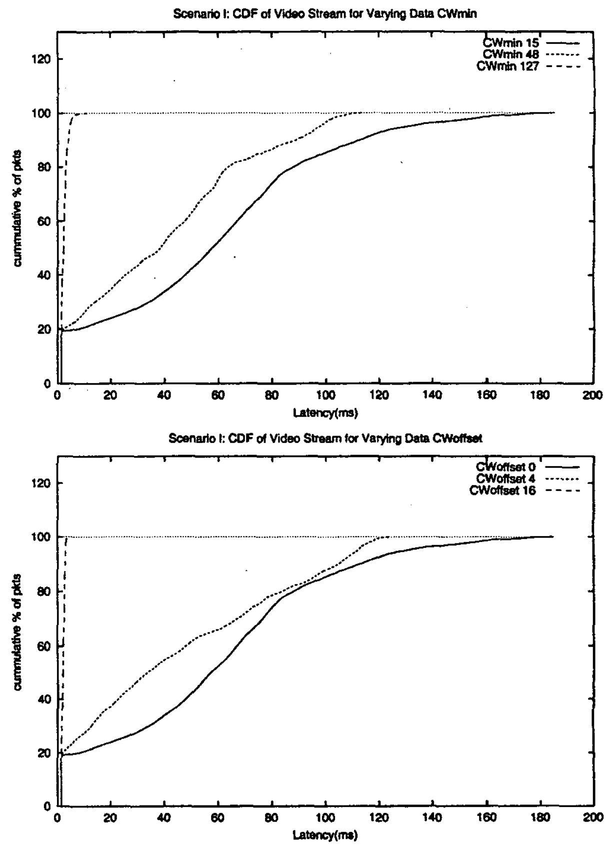

Figure 4: Tuning EDCF Parameters

6.2 Performance of HCF

Figure 5 shows a bandwidth comparison of EDCF and HCF simulated on Scenario 2. Figure 6 gives the latency CDF for DCF, EDCF and HCF.

The bandwidth for the high priority video stream drops significantly in DCF. Even the audio stream is not able to get its small required bandwidth of 0.03 Mbps; it only gets 0.02 Mbps. Also, interestingly, two of the background streams actually perform better than the audio and video while the third stream of the same traffic type is much worse (six times poorer maximum latency). This difference in behavior of the

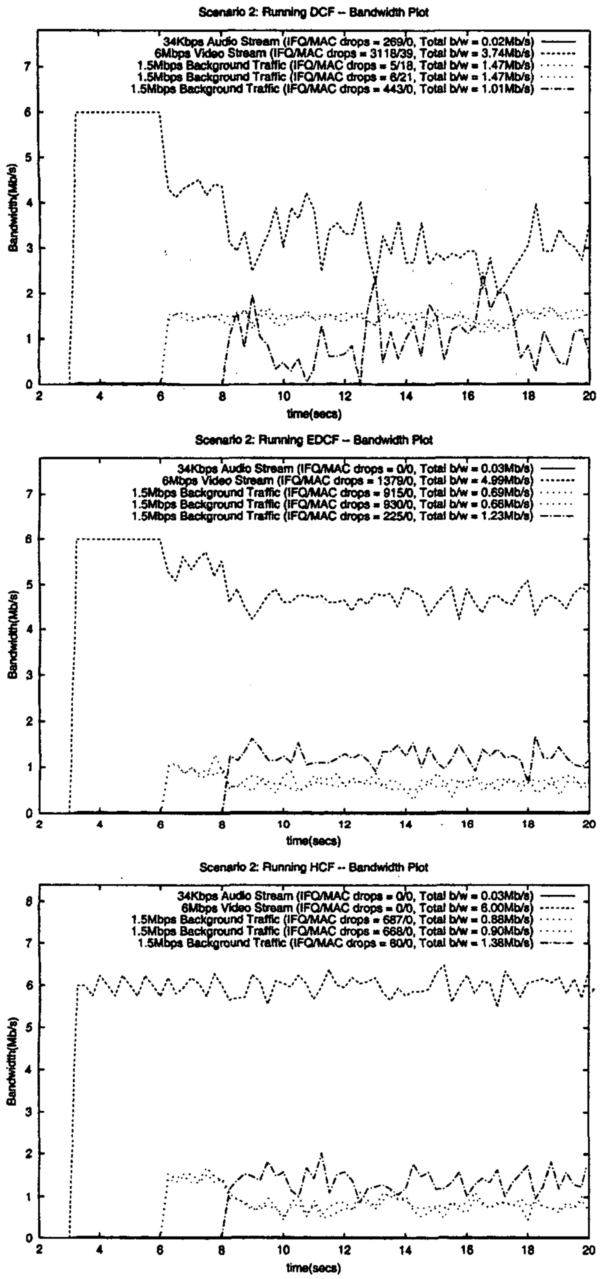

Figure 5: EDCF and HCF bandwidth usage

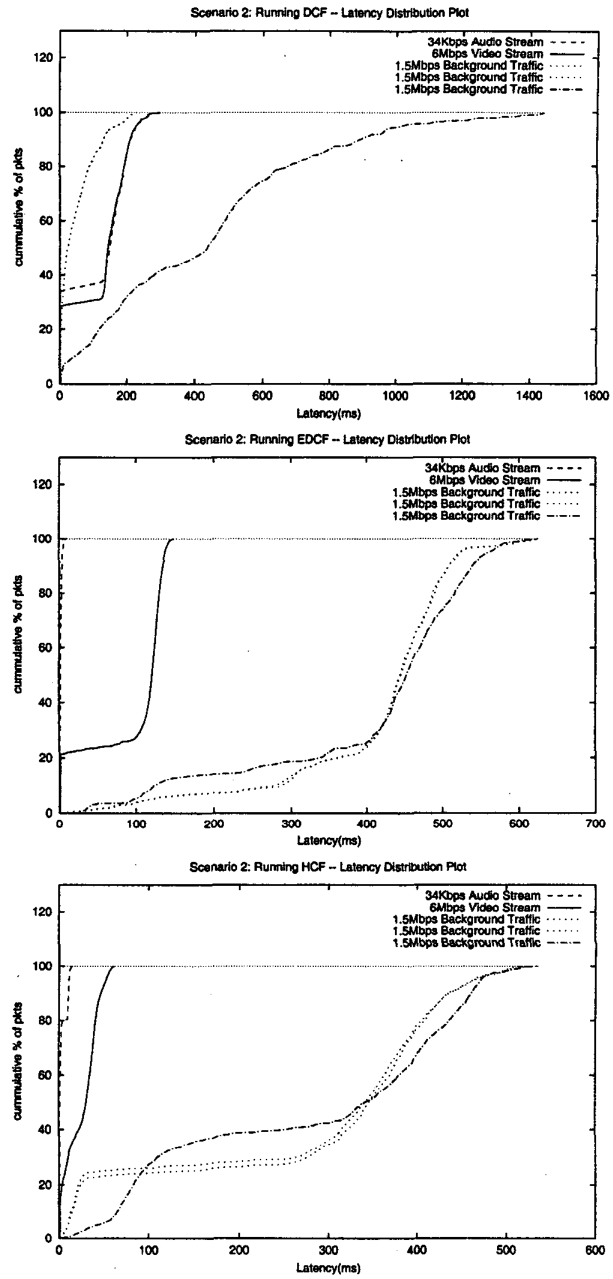

Figure 6: DCF, EDCF and HCF latency distributions

three similar background streams is an artifact of the topology, which is an undesirable characteristic.

In EDCF, once the parameters are tuned (Audio CWMin is 7 vs. a default of 15, Background Flows CWOffset is 4 vs. a default of 0), we see significant QoS improvements over DCF. Tuned EDCF is able to meet the latency constraints of the audio stream of 10 ms, which saw greater than 200 ms of latency in DCF. The video stream's worst case latency and average latency also drop, from an order of 200ms in DCF to around 100ms in EDCF, with lower jitter as well. Also, unlike DCF, the three background flows behave similarly to each other with regard to latency. However, there is a significant cost to the background flows to achieve the improvement in audio and video QoS; the net bandwidth of the background flows drops from 3.94 Mbps to 2.6 Mbps, a 33% drop, the difference being made available to the higher priority flows. Note that EDCF does not increase overall throughput of the channel (7.62 Mbps in EDCF vs. 7.71 Mbps in DCF). This is expected since the mode is fundamentally based on distributed contention just like DCF, it only adds flow-based differentiation to this contention.

In HCF as in EDCF, the audio stream is able to meet its bandwidth and latency requirements. The video stream is also able to get its bandwidth of 6 Mbps with a lower maximum and average latency of 50ms, which is much better performance than EDCF (video bandwidth of 5 Mbps with latency 120ms). Despite this increased video bandwidth, HCF gives higher bandwidth to the background flows (3.16 Mbps) than EDCF (2.6 Mbps) and improves latency as well (approximately a 15% drop). This is the result of significantly higher overall channel utilization in HCF vs. EDCF: 9.19 Mbps vs. 7.62 Mbps, about a 20% increase. We attribute this to the reduced contention overhead in HCF.

7 Future Work

Further work is needed to refine the modeling of many elements of the simulation environment of EDCF and HCF and of the protocols themselves. The channel could be improved from having uncorrelated errors to be made more representative of the errors seen in a wireless channel. We have not considered 802.11's rate adaptation, which can affect protocol behavior. Currently the simulations are for only one access point. Multiple access point scenarios create cross-cell interference particularly hurting polling schemes due to synchronization. A richer set of test scenarios should represent a wide enough variety of situations of interest to make simulation results more general. Also, we could design and evaluate various HCF scheduling algorithms.

8 Conclusion

EDCF provides significant improvements for high priority QoS traffic, however these improvements are typically provided at the cost of worse performance for lower priority traffic and the EDCF parameters can require significant tuning to achieve these performance goals. Also, EDCF does not improve channel utilization over DCF, which means it has significant overhead (about 30% is observed in our simulations.) Despite these problems, we find EDCF attractive because of its simplicity and decentralized nature.

HCF, just like its predecessor PCF, provides for much more efficient use of the medium when the medium is heavily loaded. HCF does a fairly good job of channel utilization. Due to reduced overhead, HCF can provide better QoS support for high priority streams while allocating reasonable bandwidth to lower priority streams. However, HCF involves state at the access point and is centralized, making for a less robust protocol. Furthermore, validating the effectiveness of HCF requires further study in more refined models, as described in Section 7.

Acknowledgments

We would like to thank Dr. Greg Chesson and Dr. Aman Singla of Atheros Communications for their guidance and support and Emre Kiciman of Stanford University for reviewing this paper.

References

[1] G. Chesson. An overview of quality of service issues in wireless networks. In *Stanford University Seminar*, 2001.

[2] Network Simulator Document. http://www.isi.edu/nsnam/ns. 2002.

[3] S. Mangold, S. Choi, P. May, O. Klein, G. Hiertz, and L. Stibor. Ieee 802.11e wireless lan for quality of service. 2001.

[4] IEEE 802.11 WG. *IEEE Std 802.11, 1999 edition International Standard [for] Information Technology-Telecommunications and Information Exchange between systems-Local and metropolitan networks.* IEEE, 1999.

[5] IEEE 802.11 WG. *Draft Supplement to International Standard [for] Information Technology-Telecommunications and Information Exchange between systems LAN/MAN Specific Requirements.* IEEE 802.11e/D2.0, Nov 2001.