New Prioritization Schemes for QoS Provisioning in 802.11 Wireless Networks

Konstantinos Choumas†, Thanasis Korakis ◇, Leandros Tassiulas†
† Department of Computer Engineering, University of Thessaly, Volos Greece 38221
◇ Department of Electrical and Computer Engineering, Polytechnic University, Brooklyn, NY 11201

e-mail: kohoumas@inf.uth.gr, korakis@poly.edu, leandros@inf.uth.gr

Abstract*—Due to the unreliable nature of the wireless medium, provisioning of Quality of Service (QoS) in wireless LANs is far more complicated than in wired networks. In order to address this challenge, IEEE 802.11e defines a framework for QoS support where packets are prioritized based on their traffic characteristics. In this paper, we propose two new QoS support schemes. One is based on a "user centric" approach and the other on a "packet content based" approach. The new mechanisms, in addition to the traffic itself, take into consideration the identification of the station that generates the traffic or the content of the traffic. Therefore, they use a second prioritization level on top of the one that is implemented in IEEE 802.11e. In the "user centric" approach, the mechanism defines groups of stations based on their MAC addresses and assigns different priorities to different groups. Under this classification, stations are served based on the prioritization of the group they belong. Among stations with same priority, traffic is scheduled based on the priorities given by 802.11e. On the other hand, in the case of the "packet content based" approach, the mechanism defines groups of words or phrases with their respective priority. A packet that includes words of a specific group is scheduled based on the priority that the particular content defines. The new schemes are simple yet efficient, since they are adapted to the realistic needs of today's WiFi networks. In order to evaluate the performance of these proposed schemes, we implement them using open source drivers in a Linux platform. We run experiments in a medium-size testbed. Experimentation results clearly demonstrate the performance superiority of the new schemes, as compared to the legacy IEEE 802.11e.

***Index Terms*—IEEE 802.11e, QoS, Priority**

I. Introduction

As the Internet becomes more and more popular, the trend for replacing the "last-hop" wired link with a wireless one is becoming more and more popular. This effort has not been easy so far due to several limitations that the wireless medium poses. However, as the prices of WiFi devices decrease and transmission rates increase, wireless networks gain more and more popularity.

As the number of wireless networks increases, there is a tremendous need for management of the wireless bandwidth. Due to the nature of the wireless medium, the number of wireless users that are connected simultaneously to an AP can vary a lot. Considering that all the users share the same bandwidth, it is hard to define a clear notion of quality of service guarantees. A representative example of this difficulty is the big fluctuation in the bandwidth that is assigned to a particular wireless station. It can be the whole available (if the station is the only active user associated with an AP) or can be the $1/N$ of the available, if N stations participate in the particular cell. The prioritization scheme of 802.11e fails to incorporate this parameter since it does not consider the nature of the users that are connected to an AP as long as the different needs and priorities among them.

In such an unpredictable environment where associated stations come and go dynamically, users should be prevented from using any bandwidth, no matter what is the number of stations in the network. In the premises of a company for example, employees should have higher priority on the use of the wireless network than visitors. As long as there is available bandwidth, both groups (employees and visitors) can equally share the medium. Once the network activity of the visitors is such that the QoS of the employees fall below a certain threshold, a management mechanism should be activated to prevent this. Such a mechanism must give dedicated bandwidth to the group of employees, regardless of the needs of the visitors.

On the other hand, there is traffic that different handling due to the importance of its content. *Contend based routing* [1] is a popular technology that route messages, not based to a specified destination, but based on the actual content of the message itself. In a typical application, a message is routed by opening it up and applying a set of rules to its content to determine the parties interested in its content. In the philosophy of the contend based routing, the proposed mechanism enhances the QoS provisioning of wireless networks by applying *contend based prioritization* to the forwarding traffic. The proposed scheme offers more bandwidth to packets with a payload that meets particular criteria that indicate desirable or important traffic, preventing unacceptable delay compared to the rest of the packets.

In order to address the above issues, we propose two new prioritization schemes that are based on the identification of the stations that generate the traffic or the content of the packets accordingly. In the first scheme we define groups of stations based on their MAC addresses and we assign different priorities to each of them. Among stations that have same priority, traffic is scheduled based on the priorities given by 802.11e. In the other scheme, the grouping is done based on a set of rules that are applying to the content of the frame and have been assigned higher priority. Using such schemes the wireless network can guarantee different QoS characteristics to different groups of users/packets, based on their prioritization level and their characteristics. In order to evaluate the performance of the proposed schemes we implemented them using the open source driver *MadWiFi* [2] and commercial WiFi cards. By running experiments in a medium-size testbed we show that the schemes perform efficiently in a real environment, differentiating the bandwidth allocation between different groups.

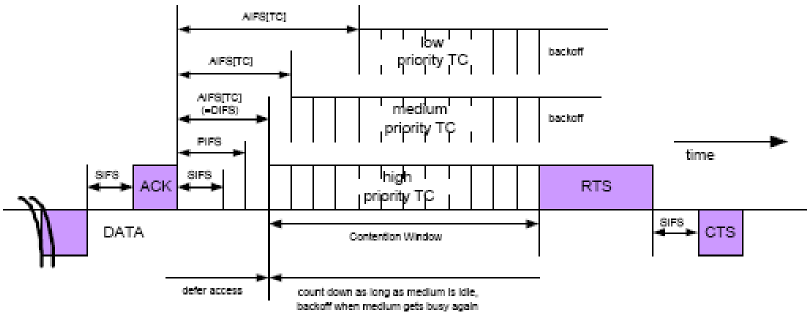

Fig. 1: Hybrid Coordination Function

The rest of the paper is organized as follows: In Section II, we give a brief description of the basic functionality of IEEE 802.11e. In Section III, we familiarize the reader with the new protocols. The implementation effort is then elaborated in Section IV. A set of measurement results along with the insights revealed therein are reported in Section V. Section VI completes the paper with final conclusions and possible future work.

II. QoS PROVISIONING IN IEEE 802.11E

The IEEE 802.11e standard [3] is an approved amendment of the 802.11 that defines enhancements of the basic MAC mechanism in order to efficiently support Quality of Service (QoS) in wireless LANs. The basic framework includes 4 priority access/class queues. These queues are used accordingly by the services of BK (Background), BE (Best Effort), VI (Video) and VO (Voice). The priority of those services increases, with the BK service having the lowest priority and the VO service having the highest one. In order for the new standard to support this prioritization, it replaces the Distributed Coordination Function (DCF) and the Point Coordination Function (PCF), with the new Hybrid Coordination Function (HCF). HCF includes two different channel access functions. The first one is the *Enhanced Distributed Channel Access (EDCA)*, that is the evolution of DCF and the second one is the *HCF Controlled Channel Access (HCCA)*, that is the evolution of PCF.

EDCF is the basic access method in IEEE 802.11e. In order to support QoS, EDCF introduces Traffic Categories (TCs). MAC Service Data Units (MSDUs) that belong to different TCs, are now delivered using different access characteristics, based on the priority of their TC. More particularly, different TCs use different inter-frame periods in order to consider the medium idle. Such periods are called *Arbitration InterFrame Space (AIFS[TC])* periods. The higher the priority of a TC, the smaller the AIFS that is used. Additionally, each TC uses its own backoff window. The *Minimum Contention Window size (CWmin[TC])* for each TC depends on the priority of the TC. The higher the priority of the TC, the smaller the CWmin is. When EDCF is active, each TC in a particular station contends for accessing the medium by starting independently a backoff procedure after detecting the channel being idle for the corresponding Arbitration InterFrame Space (AIFS). The backoff counter for each TC is a random number drawn from the interval [1, CW[TC]+1]. Similarly to the 802.11, the Contention Window (CW) follows an exponential increase every time the TC experiences a collision.

Priority	Traffic Category	Access Category	Designation
Lower	1	AC_BK	Background
	2	AC_BK	Background
	0	AC_BE	Best Effort
	3	AC_BE	Best Effort
	4	AC_VI	Video
	5	AC_VI	Video
	6	AC_VO	Voice
Highest	7	AC_VO	Voice

TABLE I: Traffic Categories

TCs correspond to the appropriate Access Categories (ACs). There is an AC for each service that was mentioned in the beginning of the section. Therefore, there is a category that is called AC_BK for the BK service, AC_BE for the BE service, AC_VI for the VI and AC_VO for the VO. TCs that correspond to the same AC have the same QoS parameters and therefore the same priority. The basic access scheme of the EDCF is depicted in Figure 1 and a table with the defined TCs is given in Table I.

An important feature of the 802.11e MAC functionality is the introduction of the *Transmission Opportunity (TXOP)*. A TXOP is a time period that is assigned to a particular station to initiate its transmissions. This period is defined by a starting time and a maximum duration. In IEEE 802.11, a station that accesses the medium has the ability to initiate a four-way handshake transmission (RTS, CTS, Data, Ack), in order to successfully transmit one data packet. In 802.11e, this scheme is extended. Once a station gets the channel, it has the ability to transmit multiple frames. The access period that is granted to this particular station is defined by the TXOP duration. The QoS parameters per TC such as AIFS[TC], CWmin[TC], TXOP(max) can be adapted over time and are announced periodically via the beacon frames.

HCCA is the second random access protocol that works as an extension of PCF. Under HCCA, the *Hybrid Coordinator (HC)*, which works as the central controller, polls stations for frame delivery. The period that the HC controls the access is called controlled contention and can be generated at any time. In order to do so, the HC requires information about the traffic needs of each station that has to be updated in a periodic basis. Based on the information about which station needs to be polled, how often, and how long a TXOP should be granted, the HC polls the stations using HCCA. The controlled contention mechanism allows stations to request the allocation of polled TXOPs by sending resource requests, without contending with other EDCF traffic. Each instance of controlled contention occurs during the controlled contention interval, which starts when the HC sends a specific control frame. This control frame forces legacy stations to set their Network Allocation Vector (NAV) until the end of the controlled contention interval, and therefore they remain silent during the controlled contention interval.

III. The proposed mechanisms

A. The "User Centric" Scheme

In the current QoS framework of the IEEE 802.11e, there is no way for the AP to identify different stations and share appropriately the medium among them. All the stations are treated equally and the QoS provisioning is done based on the traffic characteristics of the existing streams. However, this is not always fair. A typical example is this of an unlocked AP which can serve any station that is located to its coverage area. Although few of the stations (or often only one) belong to the owners of the AP, all the associated stations have the same priority and therefore, share equally the bandwidth.

The above issue can be resolved by defining priorities for different groups of associated stations. In the previous case, a fair solution would be to give higher priority to the owners of the AP and allow the other stations to share the rest of the bandwidth. As long as there is enough bandwidth for all the users, the lower priority stations will get the bandwidth they need. Once the demand of the stations exceed the available bandwidth, the higher priority stations will get the bandwidth they need, reducing the QoS for the low priority stations.

There are many ways to prioritize the associated stations. In this paper, we extend the QoS mechanism defined in 802.11e, in order to add a new priority level based on the identity of the stations. In 802.l1e each AP/station maintains different access class queues. Every packet that is ready to be transmitted is pushed to the appropriate queue. The choice of the queue is based on the service that generates the packet. Therefore, if the service is critical and requires smaller delay, the corresponding queue has higher priority. In the new scheme, those queues are used in a different way. In the first level of prioritization the AP checks the MAC address of the receiver of a particular packet. If this station belongs to the group of stations that require higher priority, the packet will be "'tagged"' as a high priority packet. In the next step, the service that the packet belongs to is examined and the packet is further classified to a higher or a lower priority queue, based on the QoS needs of the particular service. In the above description we used two levels of prioritization: high and low priority. The proposed scheme can easily extended to multiple groups of stations with different relative priority for each group.

B. The "Content-based" Scheme

In the "Content-based" scheme, the prioritization process is based on the content of the packets. Particular criteria are defined in order to examine the payload of a frame and decide about its priority. Usually, such criteria are defined using complicated queries and are expressed in XML. A simple approach of content based rules is the examination of the existence of particular words or phrases in the payload of the packet. In our scheme, packets are prioritized after their payload is examined based on such criteria. If a particular word or phrase is part of the payload, the packet gets higher priority. Otherwise the packet is treated as a normal packet and is gets the default priority. This scheme can be implemented using two levels of priorities, similarly to the first one. In the fist level, packets are classified based on their content. In the second one, packets of the same priority in the first level, get further classification by examining their traffic characteristics. Although this scheme can use several priority levels for the contend based classification, in this paper we only consider two priorities: high priority if the payload contains one of the particular phrases, and normal (low) priority for the rest of the frames.

The definition of the criteria that should be considered during the examination of the payload and the decision about the priority of the packet can significantly vary and is out of the scope of this paper. In our implementation we adopt the simple approach we mentioned earlier where we examine the existence of particular words or phrases into the payload.

IV. Implementation of the proposed schemes

For the implementation of the proposed schemes we used an open source drivers platform. More particular we modified the MadWifi [2] driver that is the Linux open source driver for commercial WiFi cards with Atheros chipsets. We chose this combination of driver-chipset (Madwifi-Atheros) because in this platform most of the MAC functionality is implemented in the driver and therefore it gives us a lot of flexibility [4]. Additionally, the particular chipset offers packet prioritization since it is has four different queues that can handle frames with different priorities. The way the packets are handled in the transmission process and how they are pushed into the four queues is controlled by the driver. Currently MadWifi has part of the 802.11e framework already implemented, defining different AIFS and maximum back-off windows for different queues [2]. A detailed presentation of the implementation of the 802.11e as a part of MadWifi is illustrated in paper [5]. In this paper the author investigates the major design requirements for SoftMAC design, and demonstrates prototypes of modifying MadWifi driver in order to satisfy a variety of requirements.

Following the proposed scheme, we modified the basic 802.11e framework of MadWifi in order to introduce two levels of prioritization: I. Prioritization based on the new classification scheme, II. Prioritization based on the service that generates the packet. In order to do so, we took the actions described below:

In the first level of prioritization, the one that is defined based on the identification of the stations or the content of the packets, we classify the frames into different priority groups. For this particular implementation in MadWifi we define two groups: A high priority group and a low priority group. However, the scheme can be easily extended to consider more than two priority groups. Based on the above, we modified the AP functionality of MadWifi and we defined two tables we call *Identity Priority Table* and *Content Priority Table*. The first table maintains information about the MAC addresses of all

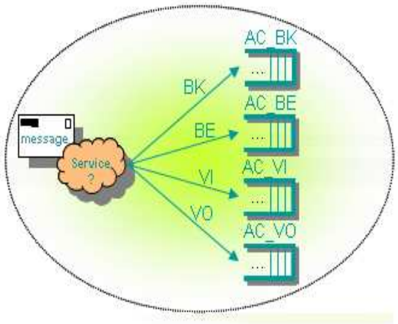

(a) 802.11e mechanism of QoS

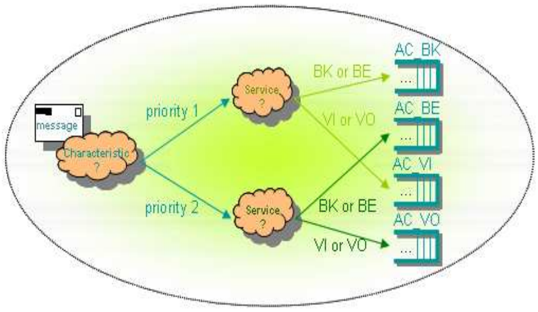

(b) New proposed mechanisms

Fig. 2: Classic and New mechanisms

the associated stations and their priority. For the "user centric" scheme we assigned priority 1 for stations of low priority and priority 2 for stations of high priority. Similarly, the second table maintains two groups of key words or phrases and their corresponding priority (again priority 1 or 2).

For the development of the Priority Tables we used a structure of MadWifi that is called Virtual Access Point (VAP). This structure is defined in each station and it keeps information about the MAC addresses of stations in the proximity and the associated AP for each of them. For the assigning of priorities to each station or phrase we developed a Graphical User Interface (GUI). The GUI, establishes a link between the driver and the user and allows the network administrator to assign different priority to different stations or phrases. A snapshot of the GUI is illustrated in Figure 3.

In order to have two levels of priority we further classified the four QoS queues that are defined in MadWifi as BK (Background), BE (Best Effort), VI (Video) and VO (Voice), into two groups (with different priorities). We assigned priority 1 (low priority) to queues BK and VI and priority 2 (high priority) to queues BE and VO.

Therefore, the MAC transmission process in the downlink of the AP has modified as follows: In the "user centric" scheme, every packet is pushed into the appropriate QoS queue, depending, firstly, on the intending receiver (and its priority) and then on the service that generates the packet.

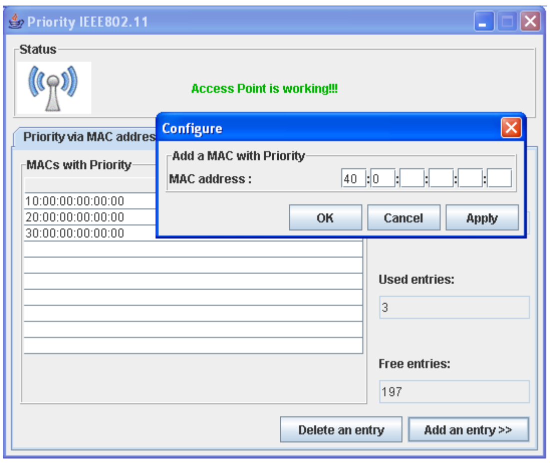

Fig. 3: GUI snapshot

In the second scheme, the criterion for giving to the packet priority 1 or 2, is based on the content of the packet. By accessing the information on the appropriate Priority Table, the driver checks whether the payload or the packet contains one of the phrases that belong to group with priority 1 or 2. If the particular packet has priority 2 (high priority), then the AP will push it to one of the two queues with priority 2, that are related with the services BE or VO. In the second stage of prioritization, if the service that generated this packet is BK or BE (lower priority services), the packet will be pushed into the BE queue. Otherwise, it will be pushed into the VO queue.

On the other hand, if the packet has low priority, it will be pushed into one of the queues BK or VI. In a case that the service that generated the packet is BK or BE, the packet will be pushed into the BK queue. Otherwise, it will be pushed into the VI queue.

The difference in the functionality between the IEEE 802.11e and the new prioritization schemes is illustrated in Figure 2. The current implementation has been done in the MAC layer of the AP and it affects the downlink traffic. In a similar way, the scheme can be extended in the client side.

V. Experiments

In order to study the efficiency of the implemented scheme, we run several experiments in real scenarios. In those experiments we considered the "user centric" scheme, that prioritizes the traffic based on the identification of the stations, we setup a testbed and we conducted experiments in a real environment. We did not repeat the experiments for the "packet content based" scheme since both schemes use the same implementation philosophy. The only difference is in the criteria of the prioritization in the first level. Therefore, after the assignation of the priority 1 or 2 to the packets, that is based on different criteria, there is no difference in the performance of the two schemes.

In the experiments we used the 802.11g mode of the cards. The topology of the first scenario, consisted of one AP and two stations. Using the GUI we described in the previous section we defined that one of the stations had high priority and one had low priority. We initiated two iperf [6] sections that generated UDP traffic. One iperf section ran between the AP and the station with high priority and the other one between the AP and the station with low priority. We should mention here that UDP traffic is considered by the driver

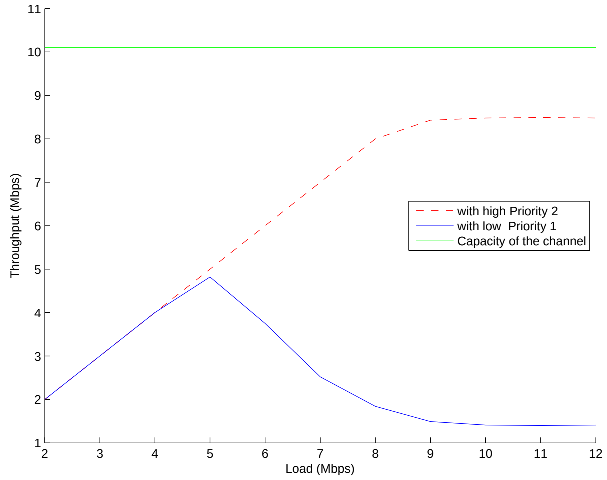

(a) Throughput of 2 stations for fixed Max Rate 12Mbps

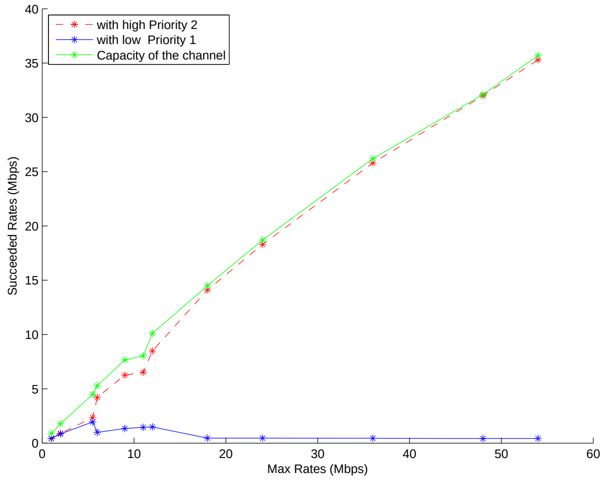

(b) Succeeded Rates of 2 stations

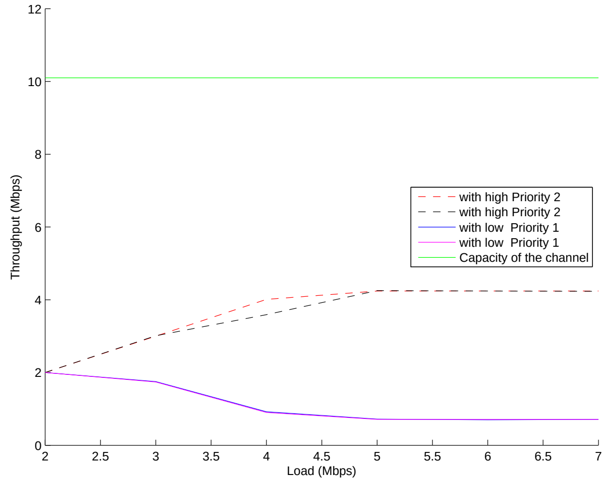

(c) Throughput of 4 stations for fixed Max Rate 12Mbps

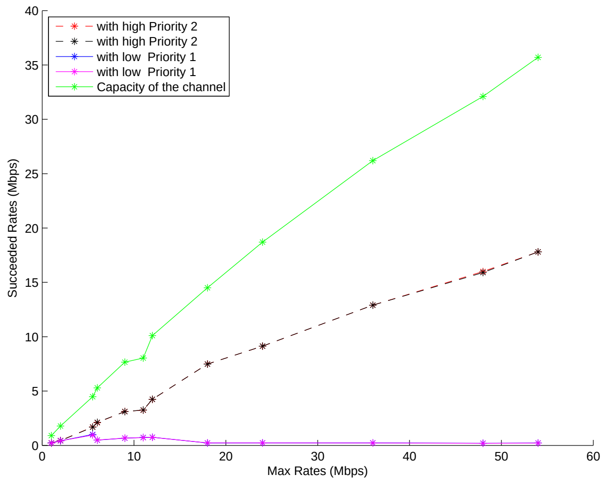

(d) Succeeded Rates of 4 stations

Fig. 4: Experimental Results

as video traffic. We ran several experiments increasing the maximum rate of the NICs in both sections from 1 Mbps to 54 Mbps. We ran each experiment for 2 min and we repeated the same experiment 5 times. We measured the average throughput in each station as the traffic load changes and for different transmission rates. In Figure 4(a) we give the results for the scenario of a fixed max rate of 12 Mbps and for different loads. As we can see in this figure, for low traffic load in the network, both stations share the same amount of bandwidth. This is because the network can sustain both the sections and therefore it provides the needed QoS to both the stations. However, as the traffic load increases, the QoS for each station changes. More particularly, once the total offered load exceeds the bandwidth of the network, the high priority station keeps receiving the needed QoS while the low priority stations starts loosing throughput. When the offered load in both the stations is high, the throughput of the low priority station is almost zero, as the high priority stations uses almost all the available bandwidth. This is compliant to the philosophy of our scheme that offers services to low priority stations only if the network recourses are not used in full by the high priority stations. Figure 4(b) illustrates the actual transmission rates for each station, as the fixed max rate increases and as the load in each station is similar to the maximum rate.

In the next set of experiments we increased the number of stations in the network. Now we have two low priority and two high priority stations. As we can see in Figure 4(c), the proposed scheme again provides equal amount of bandwidth to all the stations, until the point that the needs of the high priority stations reaches the limit of the network. After this point, the AP keeps serving the high priority stations while it shares the remain bandwidth between the low priority stations. Again in high load conditions, the high priority stations share the available bandwidth, while the low priority stations experience almost zero throughput. Figure 4(d) illustrates the corresponding successful rates for the scenario of the four stations.

Furthermore, it is worth mentioning that the scheme can easily be extended in order to incorporate a feature that would give the network administrator the ability to define the portions of the bandwidth that would be provided to the groups of each priority.

Above results are obtained in experiments that rely on large file transfer traffic patterns. In order to obtain more

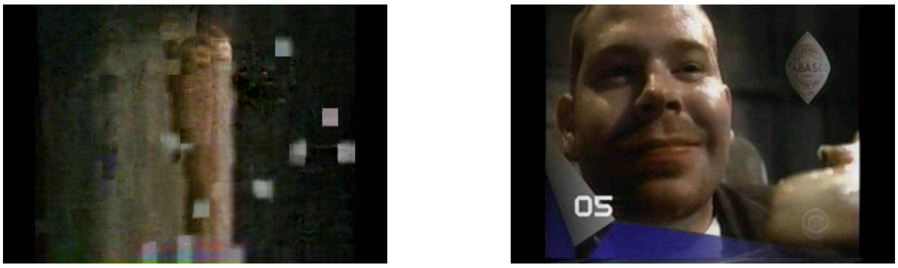

(a) Legacy 802.11e (b) new Prioritization Scheme

Fig. 5: Video Quality Comparison: A Snapshot

insights into the performance of the implemented scheme, we also considered video applications. To this end, we setup the scenario of Figure 6. The topology of this scenario consisted of one AP and three stations. One of the stations was a low priority station ($station1$). The second one was a high priority station ($station2$) and the third one generated voice traffic into the network and it did not participate in the video transmission/reception.

Two sections of video transmission were considered in the described scenario. Two VLC [7] servers were placed at the AP and were constantly streaming different commercial video clips to stations 1 and 2. The destination stations ran a VLC media player to play their video. Additionally to the traffic generated by the two video streams, an $iperf$ [6] video stream was running periodically from $station3$ to the AP, in order to increase the traffic load of the network. We alternated on-the-fly the MAC protocol in the network between 802.11e and the new prioritization scheme. We observed the changes in video quality at station 1 and 2 for the different MAC protocols.

As we can see in Figure 6, $station3$ participated in the network and generated heavy voice traffic (using a special flag in $iperf$). Since voice has higher priority than video, voice traffic kills the video traffic and therefore the video quality was poor in both video receivers. Noticeable freeze and distortion occured frequently. However, once the MAC protocol switched from 802.11e to the new prioritization scheme, the AP gave priority to $station2$ (high priority station) over any other station in the network. Therefore, the video of $station2$ was smooth and had very good quality. Figures 5(a) and 5(b) provide a snapshot of the video taken at $station2$, the first when 802.11e was active and the second one when the new scheme was active. The comparison of these two figures is typical and reveals the substantial improvement in the video quality that the new scheme can deliver.

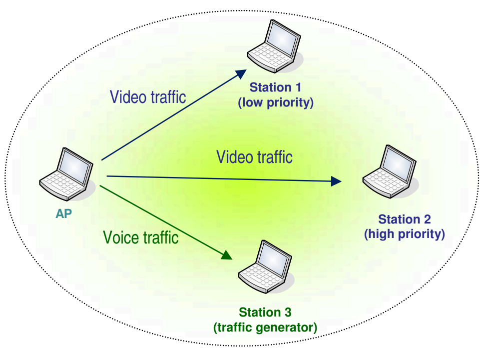

Fig. 6: The Network Setup for the Video Quality Comparison

VI. Conclusions

In this paper we proposed and implemented two new priority schemes for infrastructure 802.11 networks. The schemes are based on two levels of prioritization. In the first level the AP prioritizes the packets based on the identity of the station that receives the packet or the content of the packet. In the second one, the packet is further prioritized based on the QoS needs of the service that it belongs to. Using this approach, the network is managed in a more efficient way, providing QoS in a more realistic way. We expect that such features will play a significant role on the access mechanisms of next generation wireless networks.

In the current implementation we focus on the downlink traffic since this is the dominating factor for congestion in today's wireless networks. We are planning to extend the implementation to also incorporate the prioritization scheme on the uplink traffic. Finally, we are planning to extend the scheme that prioritize packets based on the content of the payload, to support a more advanced classification mechanism. In such a scheme, the driver will use a variety of rules that would rate the content of the packet in order to assign the appropriate priority.

References

[1] A. Carzaniga and A. L. Wolf, "Forwarding in a Content-based Network," in *Proceedings of SIGCOMM 2003*, 2003.
[2] "MADWiFi: Multiband Atheros Driver for WiFi," *http://madwifi.sourceforge.net/*.
[3] *IEEE 802.11e WG, Medium Access Control (MAC) Enhancements for Quality of Service, IEEE 802.11e Standard*. 2005 ed., 2005.
[4] T. Korakis, Z. Tao, S. Makda, and S. Panwar, "Madwifi, HostIP or Intel? A Surveillance Guidance in the Jungle of MAC Protocols' Implementation using Open Source Wireless Drivers," 2007.
[5] L. Jinsong, "SoftMAC in Heterogeneous Wireless Network," *http://hdl.handle.net/10012/3744/*, 2008.
[6] "Iperf: The TCP/UDP Bandwidth Measurement Tool," *http://dast.nlanr.net/Projects/Iperf/*, 1999-2005.
[7] "VLC Media Player," *http://www.videolan.org/vlc/*.