

Priority-based Transmission in IEEE 802.11 Networks

Tomasz Janczak and Józef Woźniak

janczak@eti.pg.gda.pl, jowoz@pg.gda.pl

Technical University of Gdańsk

Faculty of Electronics, Telecommunications and Informatics

Key words: 802.11, MAC, QoS, priorities

Abstract: Nowadays, Internet faces growing demand for carrying on multimedia traffic. Therefore, many standardization bodies work on solutions, which introduce service differentiation and QoS guarantees for Internet protocols. However, all new modifications made in the network (IP) layer would be of little value without complementary solutions implemented in lower layers. Applying of service differentiation mechanisms into the MAC (Medium Access Control) level is a challenging task, especially for wireless Local Area Networks (WLANs), due to their dynamic nature. In this paper, we propose a QoS support scheme for the IEEE 802.11 standard, based on strict frame priorities.

1. INTRODUCTION

Wireless LANs have been rapidly expanding during the last decade. The reasons are both the growing demands for cable-free communications, as well as advances in portable computers and semiconductor technology. Although, first WLAN solutions were intended as cordless replacement for Ethernet networks, now it becomes evident that WLANs must offer wider functionality, including support for multimedia traffic.

Today, two standardisation bodies lead in developing solutions for wireless LANs, namely IEEE and ETSI (European Telecommunications Standards Institute). Since the IEEE 802.11 [6] standard offering 2 Mbit/s data rate was established in 1996, there has been done much work on increasing the offered transmission speed. One task group within IEEE focused on increasing data rate using the old 2.4 GHz band. As a result, the

IEEE 802.11b supplement was published in 1999, offering two additional data rates of 5Mbit/s and 11Mbit/s. The second task group is working on a new PHY layer definition, operating in the 5 GHz band. This group has already published the IEEE P802.11a specification based on orthogonal frequency-division multiplexing (OFDM) permitting data transfer rates up to 54 Mbit/s. Although various physical layers are available, the MAC-level protocol remains the same and it does not differentiate traffic classes in its basic mode. Optionally, however, the 802.11 subnet can contain a central coordinator. The controller polls other stations, thus enabling pseudo-synchronous transmission.

While IEEE was working on 802.11, ETSI proposed the High-Performance Radio LAN (HIPERLAN/1) standard, offering up to 18 Mbit/s user bit rate [2]. Unlike in the IEEE solution, the HIPERLAN/1 MAC protocol introduces frame priorities, hence it inherently supports multimedia traffic. In spite of that advantage, HIPERLAN/1-based devices have not become widely deployed. Therefore, a newly formed ETSI working group, Broadband Radio Access Networks (BRAN) developed the HIPERLAN type 2 solution offering data rates similar to IEEE 802.11a. The new ETSI standard follows the concept of a central controller to ensure quality of service support.

Since both IEEE 802.11a and HIPERLAN/2 offer similar bandwidth capabilities, it is quite likely that other features, such as traffic class differentiation, will actually decide which protocol will gain the leading position in the market.

That is why, the IEEE 802.11 Task Group E (TGe) works intensively on introducing service differentiation mechanisms to the MAC protocol. In this paper, we briefly describe the most recent proposals from TGe and then we introduce our own solution.

1.1 IEEE 802.11 MAC Protocol

The IEEE 802.11 standard covers two lowest layers of the OSI model, namely the physical and the data link layer. In the article, we address only the MAC (Medium Access Control) sublayer of the data link layer, as it is crucial for multimedia support. The 802.11 MAC protocol can operate in two modes: DCF (Distributed Coordination Function) and PCF (Point Coordination Function).

The DCF option is the basic and mandatory operational mode, with a fully distributed channel access algorithm to the shared medium. The DCF algorithm uses the carrier sense multiple access (CSMA) protocol with collision avoidance (CA) [6], similar to that encountered in Ethernet.

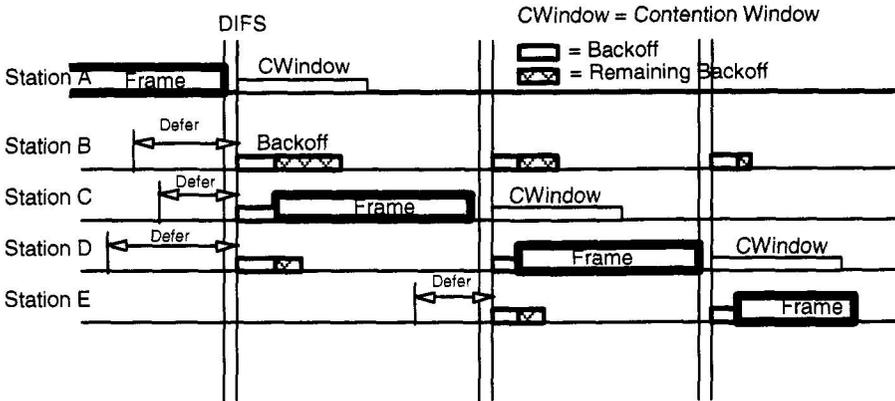


Figure 1. IEEE 802.11 Channel Access – DCF Mode

According to the DCF scheme (see Figure 1), each station wishing to transmit a frame must first sense the channel being idle for at least DIFS time (DCF Inter Frame Space). Once the channel is detected free, a station waits a random backoff period before it finally starts transmitting a frame. Similarly as in Ethernet, the random deferral minimizes collisions between multiple stations. The length of the backoff period is derived from the following formula:

$$\text{Backoff length} = \text{Random}(0, CW) * aSlotTime$$

The slot duration, $aSlotTime$, depends on the network round-trip propagation delay. The number of backoff slots is derived from a uniform distribution over the interval $[0, CW]$, where the contention window (CW) parameter ranges from a minimum value of $aCWmin$ and a maximum value of $aCWmax$. Initially, the CW parameter is set to $aCWmin$. The contention window is doubled upon each consecutive unsuccessful transmission attempt by a station, until it reaches the maximum value of $aCWmax$. CW is again reset to its minimal value, once a frame is transmitted successfully.

If a station detects any transmission during its backoff period, it suspends backoff countdown until the next access cycle.

Although the DCF mode is suitable for most data applications, it does not meet Quality of Service (QoS) requirements for multimedia applications such as voice and video transmissions. In order to support real-time traffic, IEEE 802.11 can optionally operate in the PCF mode, when a central scheduler polls other stations enabling pseudo synchronous traffic streams to be transferred. However, that centralized PCF mode is not widely deployed, and DCF remains the most popular solution. Therefore, a lot of work is

carried out to enhance DCF operation, and introduce prioritisation mechanisms for different traffic (application) types.

1.2 Related Work – Enhanced DCF Solutions

There have been several proposals on enhancements to the IEEE 802.11 Medium access control (MAC) protocol that introduce traffic class differentiation according to QoS requirements. Some early proposals include Whitecap protocol [7] and Black-Burst protocol [7]. Although both provide expedited access to network for real-time traffic, they bring significant modifications to the original algorithm. As a result, backward compatibility with the legacy solution is not achieved.

Therefore, the 802.11 Task Group E began to work on defining an enhanced DCF (EDCF) scheme, which would bring a QoS approach into DCF without major changes to 802.11 fundamentals. Very recently the TGe published three alternative draft solutions for EDCF, namely pDCF, vDCF and TCMA DCF – see Figure 2. It is expected that TGe will select one EDCF solution later on this year.

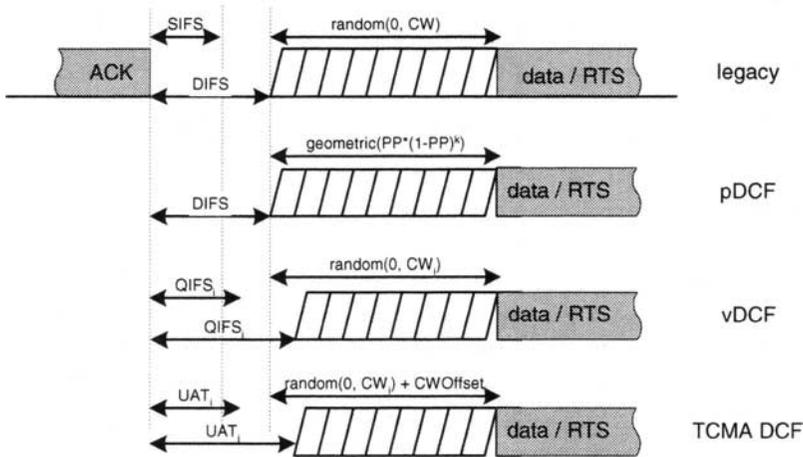


Figure 2. Enhanced DCF Proposals from IEEE 802.11 TGe

In the following sections we briefly outline the present TGe proposals, and then we describe our own solution.

1.2.1 Persistent DCF (pDCF)

The pDCF [3] scheme defines traffic different category permission probabilities (TCPPs) with which frames belonging to the traffic categories

(TCs) of the corresponding priorities are transmitted. Larger TCPP values increase the probability that frames of the corresponding TCs are transmitted. There are eight TCPPs – TCPP₀, TCPP₁, ... and TCPP₇ in accordance with eight 802.1p priorities.

Under pDCF, a station holding buffered frames of multiple TCs sums up the corresponding TCPP values to obtain a station-based permission probability (PP). The PP value denotes the probability with which the station transmits a frame after each idle slot during the backoff period:

$$PP = \sum_i TCPP_i ; \quad TCPP_i = 0 \text{ if there is no queued data of priority } i$$

This is in essence p-persistent CSMA, since effectively a station defers transmission for a *k* backoff slots given by the geometric distribution:

$$P(K = k) = PP * (1 - PP)^k$$

When the backoff timer reaches zero, the station chooses one frame from its local TC queues. First, the station draws a random number *X* in range [0, PP]. Then, it selects a frame from the local TC of priority *i* as determined by the range into which *X* falls (see Figure 3). Since the probability of transmitting a frame from its local TC of priority *i* is equal to TCPP_{*i*}, pDCF ensures fair contention among frames, according to their TCPP weights.

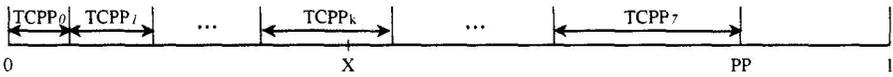


Figure 3. Local QoS Queue Selection Criterion

The pDCF scheme also provides for a special coordinator station that updates and broadcasts traffic category permission probabilities (TCPPs), adapting them to current load conditions.

1.2.2 Virtual DCF (vDCF)

Similarly to the pDCF proposal, vDCF enables prioritised access to the wireless medium for differentiated traffic categories (TCs). A vDCF station operates according to the same general rules that are defined for the legacy DCF. However, it maintains dedicated output queues per priority, and resolves contention between local queues within a station. A vDCF state machine instance controlling a virtual queue *i* is characterized by two

attributes: prioritised DIFS (called $QIFS_i$) and the lower boundary of contention window (CW_{min_i}).

QoS Interframe Space (QIFS)

In contrast to the original algorithm, vDCF state machines use different listening intervals, $QIFS_i$, to sense the channel idle condition [4]. Such concept has been also suggested in former works [1]. In vDCF, the length of the $QIFS_i$ period is set individually for each traffic category, according to the below formula:

$$QIFS_i = SIFS + aQIFS_i * aSlotTime$$

The default value for $aQIFS_i$ is 2 slots, which makes default $QIFS_i$ equal to DIFS time in the original DCF. Diversity of $QIFS$ delays in accessing channel imposes bandwidth and latency differentiation between TCs. It also ensures that low-priority traffic defers to high-priority traffic, and slower backoff counting rate is achieved for lower-priority traffic.

Backoff Calculation

A station calculates and maintains a backoff time and contention window for each priority queue i . The backoff calculation follows the DCF method, but various CW_{min_i} values are employed per traffic class, rather than one global CW_{min} . The backoff length is drawn from the following formula:

$$Backoff\ time = Random(0, CW_i) * aSlotTime$$

where CW_i is an integer within a range of aCW_{min_i} and aCW_{max} (or optionally aCW_{max_i} if available).

Discrimination of CW boundaries entails access delay differentiation between TCs. It also causes that lower-priority TCs select random backoff counters from longer CWs, thus on average receiving fewer transmission opportunities than higher-priority TCs picking from short CWs.

1.2.3 Tiered Contention Multiple Access DCF (TCMA DCF)

The TCMA proposal extends vDCF by introducing the contention window offset (aCW_{Offset_i}) for each traffic category i .

Urgency Arbitration Time (UAT)

Detection of channel idle conditions is the same as in vDCF, except that TCMA uses another terminology. Instead of $QIFS_i$, the urgency arbitration time (UAT_i) is defined for each traffic category i :

$$UAT_i = SIFS + aASC_i * aSlotTime$$

The urgency arbitration slot count (ASC_i) plays the same role as $aQIFS_i$ in vDCF, and is also equal to 2 slots for the default traffic class.

Backoff Calculation

TCMA privileges high-priority packets even more than vDCF, by adding obligatory offsets to random contention window length:

$$Backoff\ time = (Random(0, CW_i) + aCWOOffset_i) * aSlotTime$$

where $CWOOffset_i$ is specified separately for each traffic category.

Additionally, TCMA defines various contention window growing factors upon collisions, $CWPFactor_i$, per each TC (in the legacy standard the CW is always doubled) and the transmit lifetime (TFL_i) for each traffic category.

TCMA also implements fast backoff adaptation scheme, where $aCWmin_i$, $aCWOOffset_i$, and the current backoff length are scaled up or down depending on the observed channel conditions. In this way TCMA takes advantage of short-term statistical fluctuations in traffic load, especially when burst traffic is considered. The feedback necessary for adaptation is based on the number of idle slot times, the number of failed or successful transmissions and an estimate of the packet arrival rate (refer to [5] for details).

2. STRICT PRIORITY DCF (sDCF)

All enhanced DCF solutions presented in the previous section privilege higher priority frames in accessing the channel. However, they bring in some problems as well. First, these solutions do not ensure that a higher priority frame will be always transmitted before a lower priority frame. In fact, if there is a lot of low priority traffic, the high priority frame can wait a significant period of time for transmission. Consequently, delay variation for high priority packets can be reasonably large, which is not acceptable for many real-time applications.

Secondly, the common concept for all solutions is to distinguish traffic classes, by differentiating the backoff length. In general, low priority traffic uses long backoffs, while high priority frames, on average, are transmitted after shorter backoffs. Such approach implies that, for low priority traffic, the backoff length is usually longer than in the original algorithm, thus increasing the MAC protocol overhead for typical data traffic.

In this article, we propose a strict priority solution, in which higher priority frames take always precedence over the lower ones. At the given

access cycle, only stations with highest priorities at the moment can compete for transmission rights. Moreover, we preserve the same conditions in calculating the backoff length for all priorities. In this way, the low-priority traffic does not experience increased overhead.

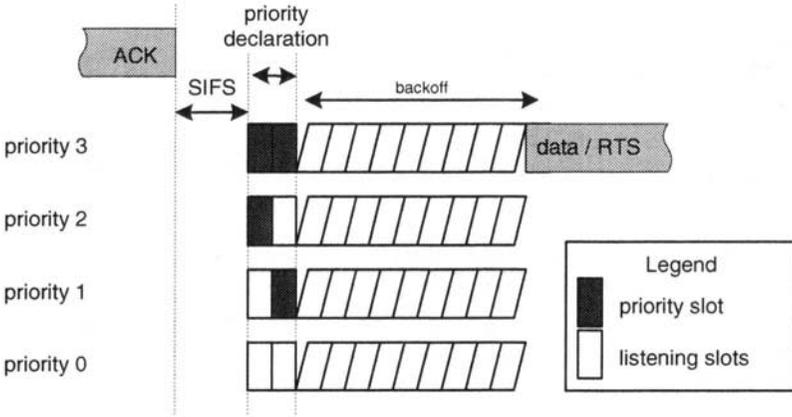


Figure 4. Strict Priority DCF (sDCF) Concept

In sDCF, each station attempting to send a frame, first transmits so-called priority declaration burst, similarly as in HIPERLAN/1 protocol [2]. The key difference is that in sDCF, the priority burst lasts always two slots, and it encodes the priority value expressed as a binary number. In other words, the priority 0 (lowest) means two empty slots, the priority 1 means one empty slot and one busy slot, etc. Figure 4 shows the possible burst combinations for all four priorities.

2.1 Example Results

In order to validate our proposal, we have examined – via simulations – the sDCF scheme behaviour. In particular, we measured data throughput and packet delay for different traffic categories.

We have taken into consideration an experimental network consisting of 20 stations, assuming fully mesh connectivity graph (no hidden stations) and 2Mbit/s channel data rate. Each station hosts a single instance of an application, which generates a stream of constant size packets (1000B long) with intervals drawn according to the Poisson distribution. To simplify simulation, all packets generated by a given source are destined to one recipient only (unicast transmission) and all streams have the same intensity. Moreover, there are 5 stations assigned to each priority class.

We have observed sDCF performance under different traffic conditions, where the offered load varies in its range from 0 up to 200 kbit/s per station. The upper bound corresponds to 4Mbit/s of total offered load by all stations.

Figure 5 shows the realized throughput per priority class, against the offered load. The dotted line shows the total throughput per station, aggregated from all traffic classes. One can see that under heavy load conditions, the total throughput stabilizes and lower-priority streams get less and less bandwidth to guarantee high QoS conditions for higher-priority streams.

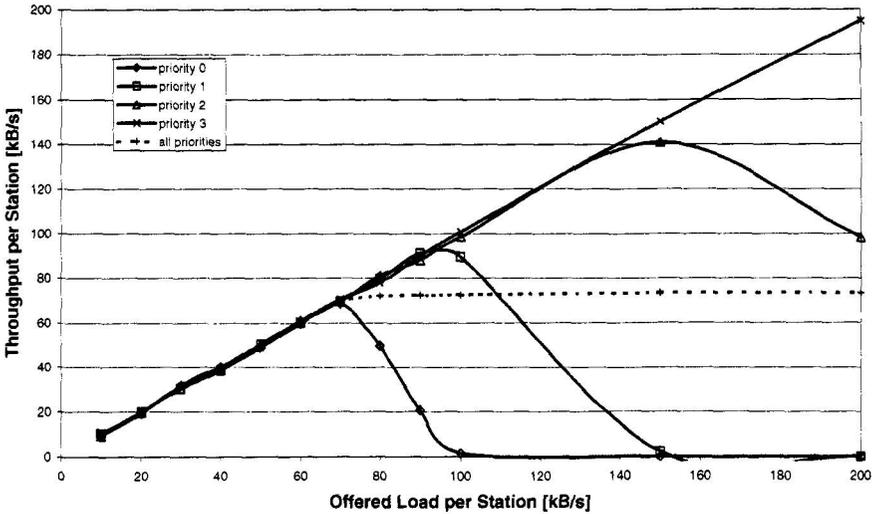


Figure 5. Data Throughput for sDCF Traffic Categories

Since lower-priority streams have less transmission opportunities when the offered load increases, the output buffers get overloaded earlier in these stations (see Figure 7). In our experiments, we assumed the output buffer capacity equal to 10kB (10 packets).

For the same reason, medium access delay in low-priority stations increases, while for high-priority stations delay remains almost unchanged under all conditions (see Figure 6).

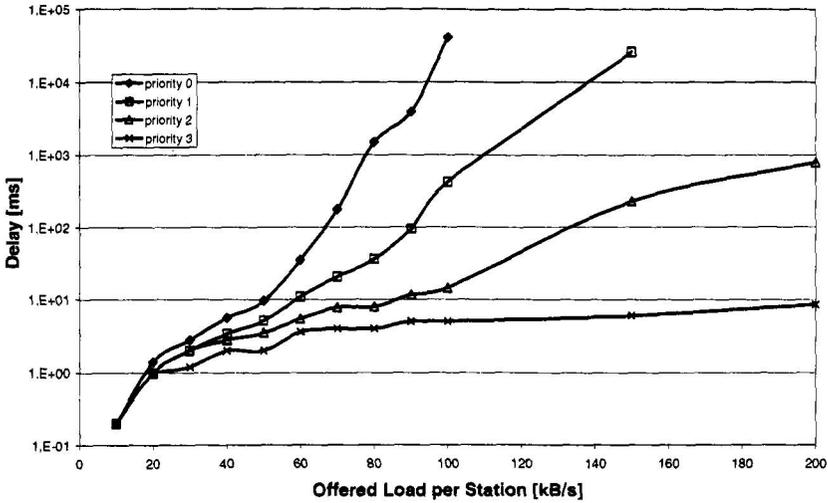


Figure 6. Packet Transmission Delay

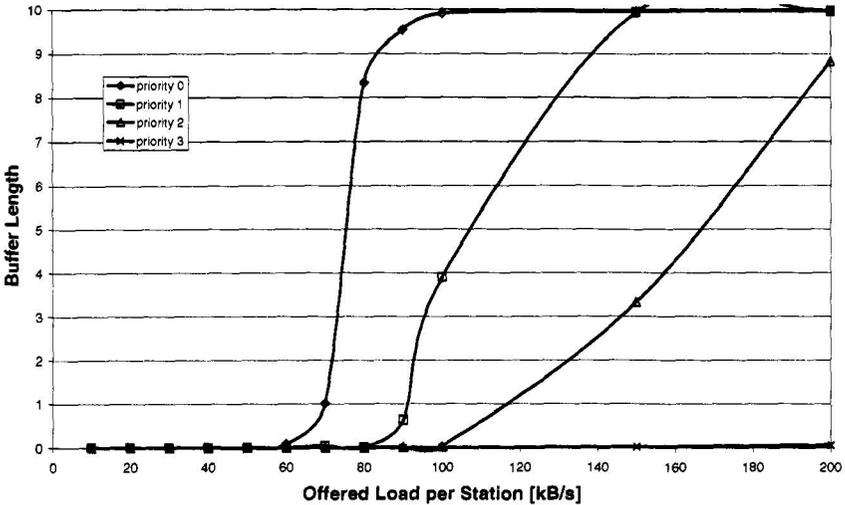


Figure 9. Buffer Length per Priority Queue

Figure 8 shows frame collision ratio observed by packets of different priorities. Under low traffic conditions, usually only one station competes for a transmission right in each access cycle. Consequently, collisions occur very rarely. When the offered load increases, output buffers in lower-priority

stations get saturated (recall Figure 7). If under such conditions no higher priorities are declared in a given access cycle, all 5 stations from the lower priority level start contending. As a result, the collision probability stabilizes at the level corresponding to 5 competing stations in each access cycle.

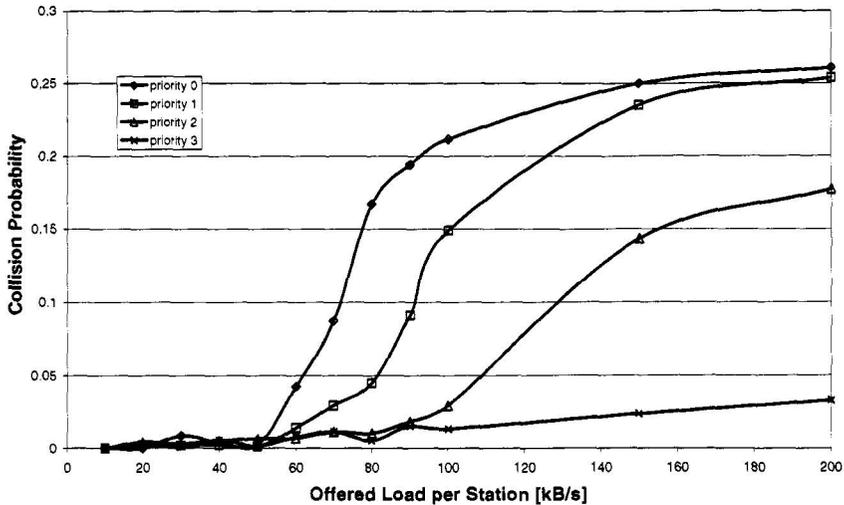


Figure 8. Frame Collisions Ratio per Priority Class

2.2 Future Work

The results presented in this paper are preliminary. They are mainly intended to verify the sDCF proposal and to address the most recent works published by 802.11 Task Group E. Future research should focus on usage scenarios involving real multimedia streams, such as voice and video applications. It is also expected, that the sDCF algorithm will be compared with other TGe proposals.

There are also several possible enhancements to sDCF, as in this paper we have not considered issues such as starvation of low-priority packets and weighted fair queuing among traffic classes. In fact, some means are needed to ensure elevating of priorities for packets that wait for transmission for a long time. Possible solutions include mechanisms based on packet lifetime, where priorities are raised as the packet buffering time approaches its limit.

Yet another extension may involve fast backoff adaptation techniques per traffic category, similar to those implemented in TCMA

3. CONCLUSIONS

In this paper we have presented the current state of work for QoS support in 802.11 standard, and proposed our own solution based on strict frame priorities (sDCF). The preliminary sDCF performance characteristics – obtained via simulation – prove that our algorithm ensures good QoS differentiation.

One can observe that both the transmission delay and the collision probability experienced by high-priority packets are relatively small (in experiments they do not exceed 10 ms and 0.03, respectively). Moreover, these characteristics remain almost unchanged under medium and heavy load conditions, making the sDCF strategy particularly suitable for supporting real-time traffic with strong QoS requirements.

At the same time the new priority-based transmission scheme does not introduce any organizational overhead to the MAC protocol. Accordingly the same level of the overall channel throughput is achieved as in the original DCF mode.

REFERENCES

- [1] Aad I., Castellucia C.: "Introducing Service Differentiation into IEEE 802.11", Proceedings of Fifth IEEE Symposium on Computers and Communications, ISCC 2000
- [2] ETSI "Broadband Radio Access Networks (BRAN); High Performance Radio Local Area Network (HIPERLAN) Type 1 ; Functional specification", July 1998
- [3] IEEE 802.11 TGe, "Proposed Text for p-DCF Contention Access Enhancement", TR-01/138 at <http://www.ieee802.org/11/>, March 2001
- [4] IEEE 802.11 TGe, "VDCF Proposed Draft Text", TR-01/131 at <http://www.ieee802.org/11/>, March 2001
- [5] IEEE 802.11 TGe, "Proposed Normative Text for TCMA with Backoff Adaptation", TR-01.135 at <http://www.ieee802.org/11/>, March 2001
- [6] IEEE Std 802.11 "Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) specifications", 1999 Edition
- [7] Sharewave, "Whitecap Network Protocols", <http://www.sharewave.com/>
- [8] Sobrinho J. and Krishnakumar A., "Quality-of-Service in Ad Hoc Carrier Sense Multiple Access Wireless Networks," IEEE Journal on Selected Areas in Communications, Vol. 17, Issue 8, pp. 1353-1368, Aug. 1999.