Paper

Modulo N Backoff Scheme for Effective QoS Differentiation and Increased Bandwidth Utilization in IEEE 802.11 Networks

Tomasz Janczak, Jerzy Konorski, Józef Woźniak, and Krzysztof Pawlikowski

Abstract—The paper presents a new modulo N channel access scheme for wireless local area networks (WLANs). The novel solution derives from the distributed coordination function (DCF) of the IEEE 802.11 standard, further elaborated as enhanced distribution channel access (EDCA) by the 802.11e draft specification. The main innovation concerns improvement of the binary exponential backoff scheme used for collision avoidance in 802.11 networks. The most appealing feature of the new modulo N backoff scheme is that it outperforms the original 802.11 solution in terms of channel utilization ratio under any traffic conditions. Furthermore, the modulo N proposal can be naturally augmented with QoS differentiation mechanisms like 802.11e extensions. The prioritized modulo N scheme achieves better throughput-delay characteristics for multimedia traffic when compared with the original 802.11e proposal. At the same time, the new solution retains backward compatibility and includes all features which have made IEEE 802.11 networks extremely popular nowadays.

***Keywords—* channel access, MAC, performance analysis, random backoff, WLAN.**

1. Introduction

Wireless local area networks (WLANs) have rapidly gained market acceptance over the last few years. The reasons are both growing demand for cable-free communications, as well as advances in portable computers and technology. Although the early WLAN solutions were merely intended as cordless replacement for Ethernet networks, it has now become evident that they must offer wider functionality, and in particular support multimedia traffic.

A significant milestone was marked by the development of the 2nd generation public wireless networks. With the emergence of 3rd generation mobile networks, broadband wireless access becomes possible. The 3rd generation systems, such as universal mobile telecommunications system (UMTS), provide enough bandwidth to support both the existing multimedia applications like speech and upcoming ones, like video-conferencing. One also observes an increased role of wirelesses networks in providing high-speed Internet access. Wireless LANs are often envisioned as a key element of 4th generation solutions for busy spots such as airports or commerce centers. In addition, the more and more popular vision of wireless homes opens up even more market opportunities for WLAN appliances. Capability of transferring high-volume multimedia streams becomes a primary goal in the design of a new generation of WLANs.

Growing demand for multimedia traffic calls for efficient bandwidth management over the scarce wireless medium. Due to the scarcity of radio resources, WLAN solutions must cope with stringent bandwidth limits, unlike their fixed counterparts. New channel access algorithms are needed to govern radio resource sharing in a way that meets multimedia application requirements while achieving high wireless medium utilization. As the speed of wireless transmission increases, the latter becomes a hot issue. At present, the medium access control (MAC)-layer protocol overhead in IEEE 802.11 networks becomes so huge that it can consume as much as 50% of available bandwidth or more [1].

This paper presents a new wireless channel access scheme, built on the basis of the IEEE 802.11 [2] and IEEE 802.11e solutions [3]. The novel proposal, called *modulo N* backoff, aims at increasing the overall utilization of a radio channel, while ensuring firm quality of service (QoS) guarantees. The novel proposal significantly outperforms 802.11e as far as the overall channel utilization ratio is concerned. Depending on traffic conditions, the modulo N backoff scheme increases the overall channel utilization ratio from 5% to 30% as compared with 802.11e enhanced distribution channel access (EDCA). The bandwidth gain depends on the number of active stations in each access cycle and the average packet size. The most appealing feature of modulo N is that under no conditions does it perform worse than its 802.11 predecessor. Furthermore, the prioritized variant of the modulo N scheme enables very effective QoS differentiation, more flexible than the original EDCA proposal from 802.11e.

The paper is organized as follows: Section 2 outlines the original channel access scheme in 802.11 networks. Section 3 introduces the concept of modulo N backoff scheme. In Section 4, optimal parameters of modulo N operation are sought. Section 5 augments the pure modulo N scheme with QoS differentiation mechanisms. Section 6 concludes the paper.

2. The IEEE 802.11 Backoff Scheme

The IEEE 802.11 standard covers the two lowest layers of the open system interconnection (OSI) model, namely the physical (PHY) and the data link layer. This paper focuses on the MAC sublayer of the latter, as it governs channel access.

The IEEE 802.11e draft [3] specification defines two operating modes for the 802.11 MAC protocol: EDCA and hybrid coordination function (HCF) controlled channel access (HCCA). EDCA is the basic and mandatory operational mode. It implements a fully distributed channel access algorithm and directly derives from the IEEE 802.11 distributed coordination function (DCF) [2]. Like its DCF predecessor, EDCA employs the carrier sense multiple access (CSMA) scheme that differs from classical Ethernet in that collision avoidance (CA) replaces collision detection.

The DCF/EDCA collision avoidance relies on the truncated binary exponential backoff (BEB) strategy, originally employed in IEEE 802.3/Ethernet networks. When an Ethernet station has a frame to transmit, it first senses the channel carrier. Once a station detects any foreign transmission, it defers until the transmission ends and then, after a fixed-duration interframe space, sends its own DATA frame. A collision occurs if two or more stations simultaneously resume transmission after deferring. Ethernet networks allow easy detection of collisions by observing changes in the signal voltage. When a transmitting station detects a collision, it delays the next transmission attempt by an integer number of slot times. The number of backoff slots is drawn from a uniform distribution from a contention window $< 0, 2^n - 1 >$, where n represents the number of the current retransmission attempt. The contention window is doubled (hence *binary* exponentiation) upon each consecutive collision, up to the predefined maximum window size (hence *truncated* BEB).

The 802.11 wireless stations implement the Ethernet backoff scheme with a modification enforced by the wireless nature of the medium. Since Ethernet-like collision detection is not possible there, IEEE 802.11 stations use a nonzero contention window from the very first transmission attempt. The contention window spans the interval $< 0, 2^{c_{\min}+n} - 1 >$, where $c_{\min} > 0$ accounts for the necessary collision avoidance in the first transmission attempt.

Once an IEEE 802.11 station senses the channel idle during DCF interframe space (DIFS) interval after previous access cycle, it defers its own DATA frame transmission for a random number of k backoff slots to minimize the probability of collisions with other senders. This scheme features a linear relationship between the random value k (backoff counter) and the number of backoff slots (Fig. 1).

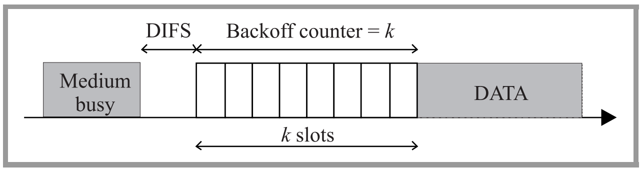

Fig. 1. Linear backoff for collision avoidance in 802.11 networks.

Though very simple and robust, the linear backoff scheme becomes a source of significant protocol overhead in wireless networks. In the IEEE 802.11a, 802.11b and 802.11g standards, the initial contention window ($CW_{\min} = 2^{c_{\min}} - 1$) has 15 slots. For example, it takes 104 μs to transmit a 512-byte packet with all MAC and PHY headers in an 802.11a [4] network operating at a 54 Mbit/s data rate. A comparable period of time is "wasted" for a single DIFS interval (34 μs) along with 7.5 backoff slots (each 9 μs long) corresponding to an average backoff time for a channel access cycle with just a single active station.

A straightforward solution aimed at reducing backoff-related overhead would be to minimize the duration of DIFS and backoff slots. Unfortunately, these time constants cannot be decreased at will, since they are determined by the propagation delay and receiver/transmitter switchover time.

Another option is to minimize the number of backoff slots. This can be achieved by:

- adapting the contention window range and/or size to current traffic conditions, or
- changing the way the backoff counter value is encoded and communicated to other stations.

While the former approach has been extensively studied (see, e.g., [5], [6]), there is little work concerning backoff coding schemes. This paper fills this gap by describing a new backoff coding scheme, called modulo N. It is specifically intended for radio environments such as IEEE 802.11 networks, where it can significantly reduce backoff overhead.

3. Modulo N Backoff Scheme

An optimal backoff algorithm should have the following properties:

- low best-case backoff length to take advantage of light-load traffic conditions;
- small average backoff length for typical multi-station channel access scenarios;
- robustness in the sense of keeping a moderate frame collision rate under heavy-load traffic conditions.

The modulo N scheme satisfies all the above requirements. It features the best case close to DCF/enhanced DCF (EDCF), but at the same time significantly improves the worst case. It also achieves a reduced average backoff overhead.

A wireless node supporting the modulo N scheme follows the BEB strategy in order to reduce the risk of DATA frame collisions, like it does under DCF. However, it employs a different backoff coding to inform other stations about the backoff counter value it has selected at random. Figure 2 illustrates the principle of modulo N operation. When a station has a DATA frame ready for transmission and senses the medium busy, it selects a random backoff counter value k. This value is next divided modulo N into an integer part $\boldsymbol{k}_{/N}$ and a remainder part $\boldsymbol{k}_{\%N}$, so that

$$k = k_{/N} \cdot N + k_{\%N} .$$

After the previous channel access cycle is finished (e.g., after DIFS, like in IEEE 802.11), a station senses the medium for the duration of $\boldsymbol{k}_{/N}$ slots. If it remains idle, a station broadcasts a one-slot *busy signal*. Next, it waits for $\boldsymbol{k}_{\%N}$ idle slots before it finally commences a DATA frame transmission. If the station detects any foreign signal during the idle slots, it is inhibited from transmission, i.e., gives up and waits until the next access cycle. This may only happen if another station has won the contention by selecting a shorter backoff. Like in the original DCF, the inhibited station decrements its backoff counter by the number of elapsed slots.

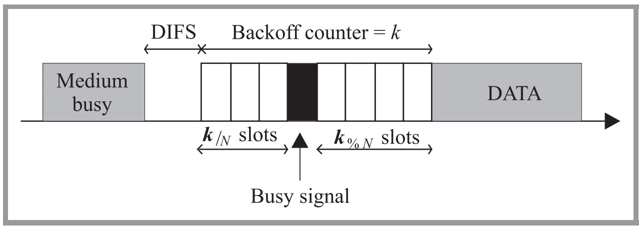

Fig. 2. Modulo N backoff encoding.

In the original IEEE 802.11 standard, the number of backoff slots is always equal to the backoff counter value. In contrast, in the modulo N scheme, the required number of backoff slots for a given backoff counter value k can be expressed as $\boldsymbol{k}_{/N} + \boldsymbol{k}_{\%N} + 1$.

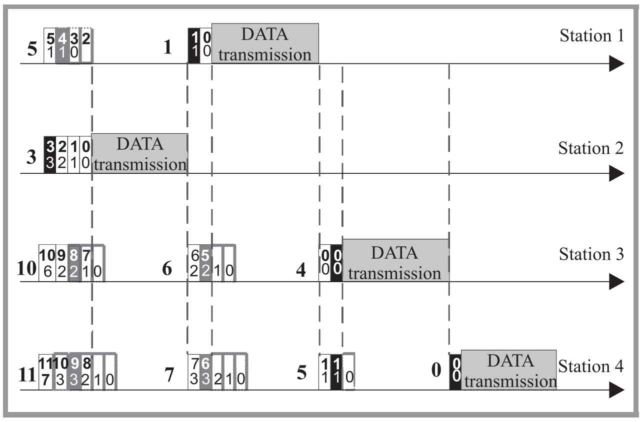

Fig. 3. Example network operation with modulo 4 backoff.

Figure 3 depicts an example network operation with four wireless stations using the modulo 4 backoff scheme (a semiformal specification of the proposed channel access algorithm can be found later, cf. Fig. 10). There are three types of backoff slots distinguished in Fig. 3:

- white boxes indicate idle slots (a station is listening);
- black boxes represent busy signal slots;
- gray boxes signify slots that would carry busy signals had a station not been inhibited by a foreign transmission.

The numbers shown on the diagram represent the value of a backoff counter:

- the numbers to the left of the slot boxes correspond to backoff counter values at the beginning of a new channel access cycle (i.e., after DIFS);
- the upper numbers inside the boxes represent current of backoff counter values at the end of a slot;
- the lower numbers inside the boxes represent theoretical of backoff counter values had a station not been inhibited by a foreign transmission; note that the upper and lower numbers are equal in a winning station.

Station 2 has the lowest starting backoff counter and it wins the first access cycle. This station starts with a one-slot busy signal, since its initial backoff counter 3 divided modulo 4 gives $\boldsymbol{k}_{/N} = 0$. All the other stations start with listening slots, and they are inhibited by the busy signal. In next three slots all stations proceed according to the original DCF algorithm. They decrement their backoff counters by one in every slot before station 2 finally commences transmission. The transmission phase includes also acknowledge (ACK) and interframe spaces. Such a transmission phase is considered a slot in DCF/EDCA, and the remaining stations decrement their backoff counters at the end of the transmission phase.

The second access cycle starts with an immediate busy signal from station 1. Again, the busy signal inhibits stations 3 and 4, which still have their backoff counters higher than 4. Note that no station decrements its backoff counter during the busy signal slot, hence station 1 listens for one more of the $\boldsymbol{k}_{\%N}$ slots before it zeroes its backoff counter.

Both stations 3 and 4 begin the third access cycle with one of the $\boldsymbol{k}_{/N}$ listening slots. As they do not detect any transmission during this slot, they both decrease their backoff counters by N (here, $N = 4$). Next, they both have the backoff counter lower than N so they announce transition to the $\boldsymbol{k}_{\%N}$ slots by sending a one-slot busy signal. Station 3 has $\boldsymbol{k}_{\%N} = 0$ (backoff value 4 modulo 4 gives a remainder of 0), and it starts data transmission immediately after the busy signal slot. Station 4 listens for one of the $\boldsymbol{k}_{\%N}$ slots and detects signal from station 3. This inhibits station 4 from commencing its own transmission.

Station 4 enters the last access cycle with the backoff counter equal to 0. Even in such a case a station has to send the busy signal. If station 4 did not send the busy signal first, its transmission could collide with a busy signal from any other station that has a backoff counter less than 4.

By examining the lower numbers inside the boxes, one can compare modulo 4 with the original DCF/EDCA backoff scheme. For example, station 1 would need only 3 slots to announce the backoff counter value of 5, and station 4 would need 6 slots to announce value of 11. On the other hand station 2 needs 4 slots with backoff counter 3, and station 4 tranmits a one-slot busy signal even if its backoff counter is equal to zero. The next section provides more detailed analysis of the modulo N scheme and its comparison with DCF/EDCA schemes.

4. Optimal Modulo N Parameters

The modulo N scheme can be subject to numerous parameterizations. The value N itself is an apparent parameter to manipulate. As N increases, the maximum backoff length decreases, but the impact upon the average backoff is not obvious given that a large N leaves little room for collision avoidance based on the integer parts of the backoff counter values, especially when multiple stations compete in successive access cycles.

Figure 4 compares modulo 4 and DCF backoff schemes assuming default IEEE 802.11a settings: slot time = 9 μs, DIFS = 34 μs, $CW_{\min} = 15$, $CW_{\max} = 1023$, and 54 Mbit/s data rate. The simulation results are provided for an ideal radio channel with all the stations within each other's range (no hidden terminals), and DATA frame errors occur only due to collisions, transmission errors being negligible.

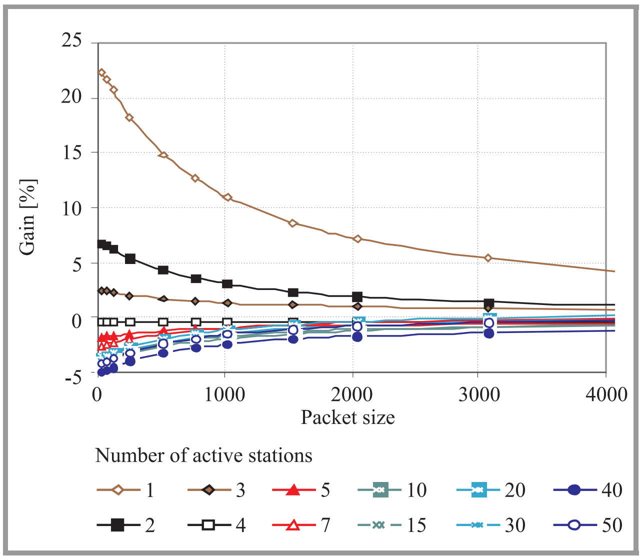

Fig. 4. Modulo $4|2^6$ throughput gain over legacy EDCA.

The curves represent simulation runs for various numbers of active stations under saturation conditions (i.e., each station always has a DATA frame ready for transmission). Let S_x denote the saturation throughput for a backoff scheme x. The percentage gain shown on the y-axis is defined as

$$\text{gain} = \frac{S_{\text{mod N}} - S_{DCF}}{S_{DCF}} .$$

From Fig. 4 it follows that the modulo 4 scheme outperforms DCF if fewer than four stations contend for channel access at one time. The highest gain is achieved in the case where only one station is active in each access cycle. This can be easily explained: with one contending station and $CW_{\min} = 15$, an average DCF backoff is 7.5 slots whereas an average modulo 4 backoff is only 4 slots.

Under extremely heavy load, with more than 20 active stations in each access cycle, an average DCF backoff becomes less than one slot, which is the lower bound for modulo N. Therefore, DCF performs better than pure modulo N under extreme traffic conditions.

An appealing feature of modulo N is that the maximum backoff window is bounded by $(CW_{\max}/N) + N$, which is a significant improvement over $CW_{\max}$ in EDCA. Assuming $N = 4$, the maximum backoff time is 260 slots in modulo 4 as compared with 1023 slots in EDCA. It makes sense, therefore, to manipulate other protocol parameters. In IEEE 802.11 networks, the contention window ranges between $CW_{\min}$ and $CW_{\max}$, which are interrelated as follows:

$$CW_{\max} = (c_{\text{inc}})^{c_{\max}} \cdot (CW_{\min} + 1) - 1 ,$$

where the IEEE 802.11a defaults are: $CW_{\min} = 15$, $c_{\text{inc}} = 2$, $c_{\max} = 6$, and $CW_{\max} = 1023$. Thus, in order to get the maximum backoff duration of 1023 slots, one could configure a modulo 4 scheme with $c_{\text{inc}} = c_{\max} = 4$. Hereafter this combination will be denoted *modulo* $N|c_{\text{inc}}^{c_{\max}}$.

As illustrated in Figs. 5, 6, and 7, higher N/c_{inc} values generally lead to better channel utilization. Unfortunately, a serious drawback of modulo $5|5^5$ is the maximum contention window $CW_{\max}$, reaching $16 \cdot 5^6$, or 250 000. Considering that the backoff length is 5 times shorter, it is still 50 000 backoff slots in the worst case. Even though

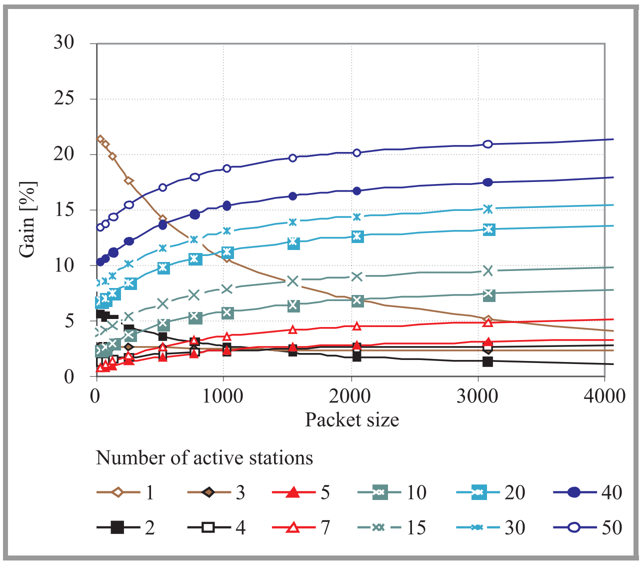

Fig. 5. Modulo $3|3^6$ throughput gain over legacy EDCA.

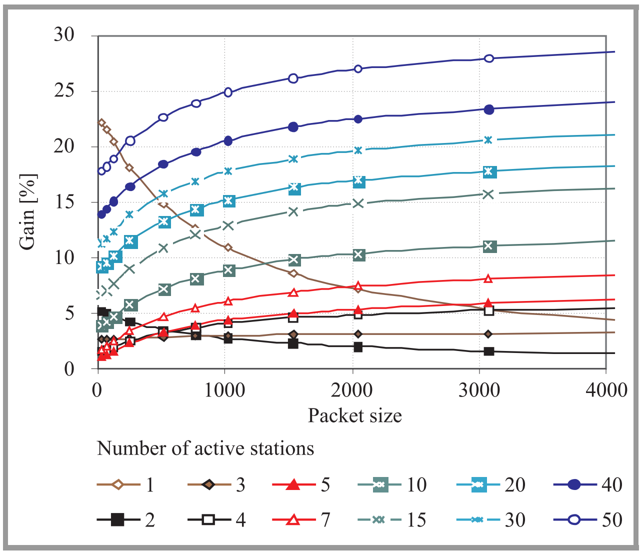

Fig. 6. Modulo $4|4^6$ throughput gain over EDCA.

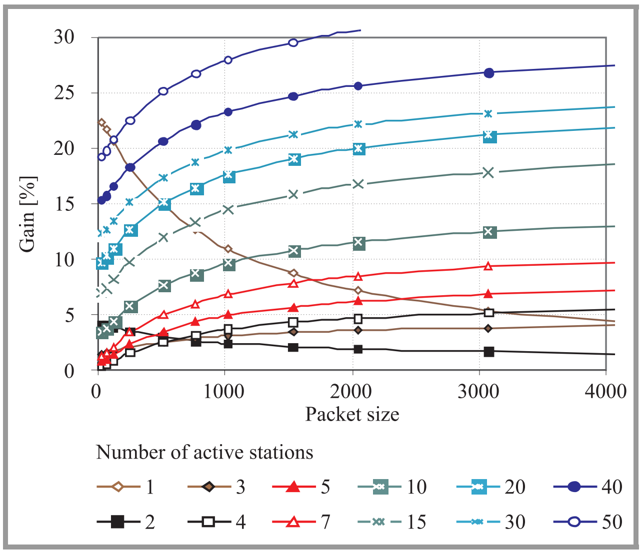

Fig. 7. Modulo $5|5^5$ throughput gain over EDCA.

the probability of reaching this limit is very small, such long-tailed distributions should be avoided in real networks. Yet similar performance results can be achieved under modulo $5|5^4$, where the maximum backoff length is limited to 2000 slots, comparable with that under DCF. This limit can be reduced even more under modulo $4|4^4$, which ensures maximum backoff length of 1024 slots, very much like under DCF.

Modulo $4|4^4$ seems a reasonable configuration choice, bearing in mind that predictable network operation is more valuable than fine-tuning the protocol parameters under very heavy load (with over 10 stations competing in each access cycle).

From Fig. 8 it can be seen that modulo $4|4^4$ configuration outperforms EDCA for all traffic scenarios. The explanation is that modulo 4 allows broadening the contention window beyond the $CW_{\max}$ limit defined for 802.11a, while retaining the maximum backoff duration of 1023 slots. Clearly, the increased *CW* range results in fewer collisions; as a consequence, better channel utilization can be achieved.

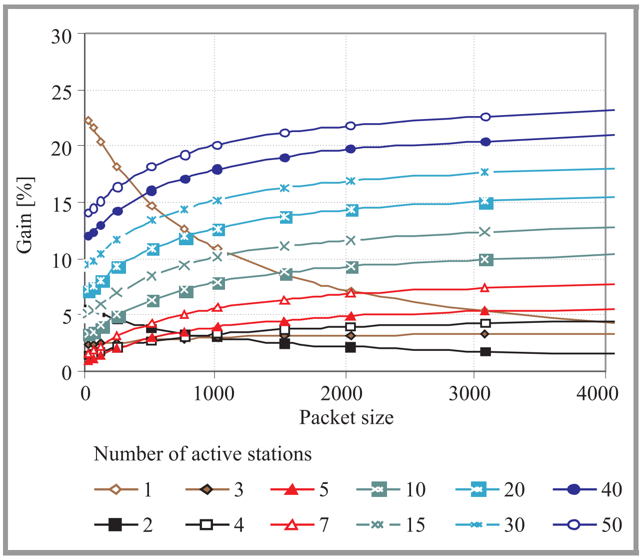

Fig. 8. Modulo $4|4^4$ throughput gain over EDCA.

The lowest gain is achieved in a scenario when two stations are active in each access cycle. This nicely tones in with existing research reports, which indicate that default DCF/EDCA parameters are optimal for two-station scenarios [5]. Notably, even in such a case modulo N performs better than the original DCF/EDCA.

5. Proritized Modulo N Backoff

The pure modulo N scheme does not allow for prioritization of traffic streams. Nevertheless, the scheme can be easily augmented with QoS differentiation as illustrated in Fig. 9.

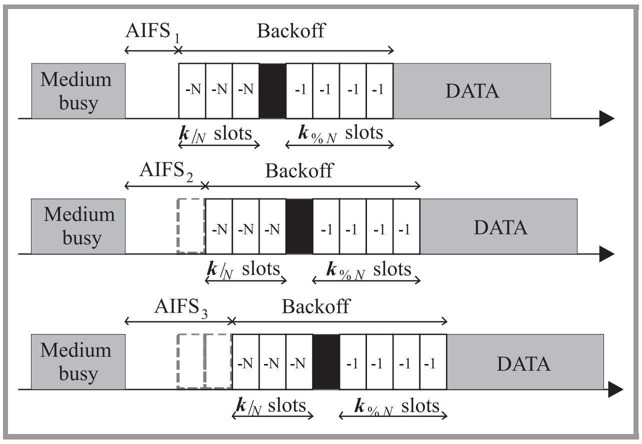

Fig. 9. Prioritized modulo N backoff encoding concept.

Like in IEEE 802.11e EDCA, the basic idea is to replace the DIFS interval with the arbitration interframe space (AIFS) intervals defined on a per-class basis, as well as to use per-class contention window ranges.

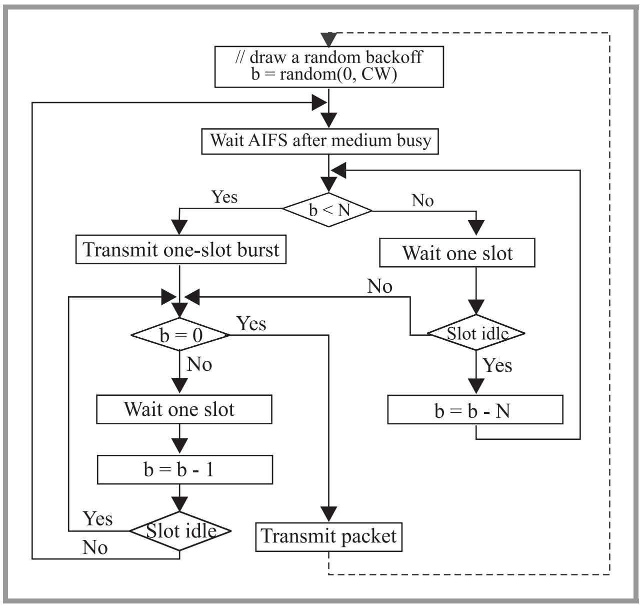

Fig. 10. Prioritized modulo N channel access algorithm.

Figure 10 gives a flowchart description of the prioritized modulo N channel access scheme. Like in EDCA, a station decrements its backoff counter both at the end of an idle slot as well as during a foreign transmission period (i.e., a DATA frame exchange sequence started by another station). Unlike in EDCA, however, decrementing should take place at the beginning of a foreign transmission period (i.e., one slot after transmission starts), and not when AIFS expires as described in the IEEE 802.11e draft. Furthermore, DATA frame transmission should be commenced immediately after the backoff counter reaches 0 (in IEEE 802.11e, a station starts transmission one slot later). In that sense, the prioritized modulo N resembles more IEEE 802.11e draft version 4.2 than the more recent version 8.0. Nonetheless, both approaches are functionally equivalent and one described below facilitates a simple implementation of the modulo N scheme.

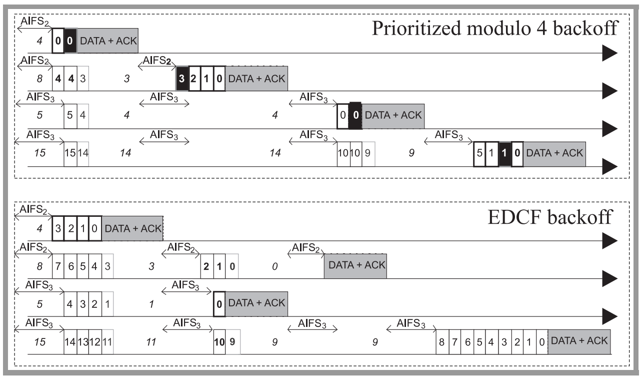

Fig. 11. Example network operation with prioritized modulo N backoff.

Figure 11 shows an example network scenario with four wireless stations using the prioritized modulo 4 and an original EDCA backoff. White and black boxes indicate idle and busy slots, respectively. The numbers to the left of the slot boxes show backoff counters at the beginning of an access cycle. Those inside the boxes correspond to the current (end of slot for modulo N and start of slot for EDCA) backoff values.

The prioritized modulo N provides more stringent AIFS-based QoS differentiation as compared with original EDCA. The reason is that in modulo N a one-slot difference in AIFS intervals corresponds to an N-slot difference in backoff counters. Consider high-priority stations (AIFSN = 2) and low-priority stations (AIFSN = 3) depicted in Fig. 11. In the very first DATA frame exchange, a high-priority station 2 decrements its backoff counter by 5, while low-priority stations 3 and 4 decrement their backoff counters just by one. In the original EDCA, the low-priority stations would decrement their backoff counters by 4, which indeed is not much less than 5 in a high-priority station.

In general, a high-priority EDCA station is unaffected by low-priority stations only if its backoff counter is already equal to 0 when AIFS expires. In contrast, transmission from a high-priority modulo N station does not depend on the presence of low-priority stations for all backoff counters less than *N*. Consider for instance the second DATA frame exchange in Fig. 11. The high-priority station 2 has a backoff counter equal to 3, which is enough to prevent low-priority stations from transmission in this access cycle (even though they had lower backoff counters initially). Similar behavior is not possible in the original EDCA. Consider again the second DATA frame exchange under EDCA. We see that a high-priority station (backoff = 3) loses in competition with a low-priority station (backoff = 1).

6. Conclusions

The paper describes the new modulo N backoff scheme. Both a semiformal description of the new channel access scheme and simulation results that compare the new scheme with existing ones like IEEE 803.11 DCF and EDCA, have been presented.

The description of modulo N reveals that its complexity is comparable with legacy schemes. At the same time, the obtained performance results show that the new scheme increases the overall channel utilization between 5% and 30% as compared with IEEE 802.11 DCF.

Furthermore, the paper describes the prioritized variant of the modulo N scheme. This variant enables very effective QoS differentiation, which is also more flexible than the original EDCA scheme of IEEE 802.11e.

References

[1] A. Doufexi, S. Armour, M. Butler, A. Nix, D. Bull, J. McGeehan, and P. Karlsson, "A comparison of the HIPERLAN/2 and IEEE 802.11a wireless LAN standards", *IEEE Commun. Mag.*, vol. 40, no. 5, pp. 172–180, 2002.

[2] "IEEE Standard for Information Technology – Telecommunications and Information Exchange between System – Local and Metropolitan Area Networks – Specific Requirements" – Part 11: "Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specification", IEEE 802.11-1999 [Online]. Available: www.ieee.org

[3] "Draft Supplement 8.0 to Standard for Telecommunications and Information Exchange between Systems – LAN/MAN Specific Requirements" – Part 11: "Wireless Medium Access Control (MAC) and Physical Layer (PHY) Specifications: Medium Access Control (MAC) Enhancements for Quality of Service (QoS)", IEEE Draft 802.11e-2005.

[4] "Supplement to IEEE Standard for Information Technology Telecommunications and Information Exchange between System – LAN/MAN Specific Requirements" – Part 11: "Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications: High-Speed Physical Layer in the 5 GHZ Band", IEEE 802.11a-1999 [Online]. Available: www.ieee.org

[5] F. Cali, M. Conti, and E. Gregori, "Dynamic tuning of the IEEE 802.11 protocol to achieve a theoretical throughput limit", *IEEE/ACM Trans. Netw.*, vol. 8, no. 6, pp. 785–799, 2002.

[6] M. Natkaniec and A. Pach, "An analysis of modified backoff mechanism in IEEE 802.11 networks", in *Proc. Polish-German Teletraf. Symp. PGTS'2000*, Dresden, Germany, 2000

Tomasz Janczak received the M.Sc. degree in computer science in Gdańsk University of Technology (GUT), Poland, in 1999. In 2001, he also received the M.Sc. degree in management and economics from GUT. Since 1999, he has been working in GUT towards the Ph.D. degree which he received in 2004. His research work focused on wireless local area networks, including topics covered by IEEE 802.11, HIPERLAN and Bluetooth solutions. He has published several technical papers on supporting quality of service and service fairness in wireless networks with dynamically changing topologies. He is a senior software architect at Intel Technology Poland. At present, he works as a system architect at Intel Corporation, at the R&D networking site located in Gdańsk.

e-mail: tomasz.janczak@intel.com
Intel Technology Poland
Słowackiego st 173
80-298 Gdańsk, Poland

Józef Woźniak received his Ph.D. and D.Sc. degrees in telecommunications from the Gdańsk University of Technology, Poland, in 1976 and 1991, respectively. At present he is a Full Professor in the Faculty of Electronics, Telecommunications and Computer Science at the Gdańsk University of Technology. He is the author or co-author of more than 200 journal and conference papers. He has also co-authored 4 books on data communications, computer networks and communication protocols. In the past he participated in research and teaching activities at Politecnico di Milano, Vrije Universiteit Brussel and Aalborg University, Denmark. In 2006 he was Visiting Erskine Fellow at the Canterbury University in Christchurch, New Zealand. He has served in technical committees of numerous national and international conferences, chairing or co-chairing several of them. He is a member of IEEE and IFIP, being the Vice Chair of the WG 6.8 (Wireless Communications Group) IFIP TC6 and the Chair of an IEEE Computer Society Chapter (at the Gdańsk University of Technology). His current research interests include modeling and performance evaluation of communication systems with the special interest in wireless and mobile networks.

e-mail: jowoz@eti.pg.gda.pl
Gdańsk University of Technology
G. Narutowicza st 11/12
80-952 Gdańsk, Poland

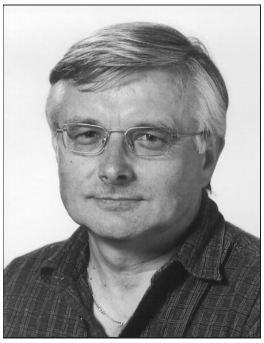

Krzysztof Pawlikowski received a Ph.D. degree in computer engineering from the Gdańsk University of Technology, Poland, and worked at that university until Feb. 1983. At present he is a Professor of computer science at the University of Canterbury, in Christchurch, New Zealand. He is the author of over 150 journal and conference papers and four books, has given invited lectures at over 80 universities and research institutes in Asia, Australia, Europe and North America. He is the principal developer of Akaroa2, a universal controller of quantitative stochastic simulation, which is used at universities in over 70 countries around the world. He was the Alexander-von-Humboldt Research Fellow (Germany) in 1983-84 and 1999, and a Visiting Professor at universities in Australia, Italy, Germany and the USA. His research interests include multimedia telecommunication networks and their performance modeling, as well as methodology of discrete-event computer simulation and distributed processing. He a is senior member of the IEEE, member of the ACM and the Royal Society of New Zealand.

e-mail: K.Pawlikowski@ieee.org
Network Research Laboratory: Protocols,
Distributed Processing and Simulation
Department of Computer Science
and Software Engineering
University of Canterbury
Private Bag 4800
Christchurch, New Zealand

Jerzy Konorski – for biography, see this issue, p. 40.