Sliding Contention Window (SCW): Towards Backoff Range-Based Service Differentiation over IEEE 802.11 Wireless LAN Networks

Abdelhamid Nafaa, Adlen Ksentini, and Ahmed Mehaoua
University of Versailles — PRiSM Laboratory
Brent Ishibashi, Youssef Iraqi, and Raouf Boutaba, University of Waterloo

Abstract

A number of works have tried to adjust the contention window in order to provide differentiated quality of service in IEEE 802.11-based wireless networks. By giving different service classes different CWs, the distribution of backoff intervals (chosen randomly, on the interval [0, CW]) will reflect the desired service classes. However, these protocols cannot deliver firm service guarantees while maintaining high network utilization, particularly under congested network conditions. In this article we propose a new MAC protocol featuring a sliding CW (SCW) for each network flow. The SCW dynamically adjusts to changing network conditions, but remains within a per-class predefined range in order to maintain a separation between different service classes. Each flow's SCW reacts based on the degree to which class-defined QoS metrics are satisfied. Simulation results show that compared to the enhanced distributed coordination function (EDCF) scheme of 802.11e, SCW consistently excels, in terms of network utilization, strict service separation, and service-level fairness.

With the improvement in bandwidth and throughput of wireless technologies, it is becoming possible to support true multimedia applications, including data, voice, and video traffic. However, each of these applications has requirements for the delivery of their packets. It is insufficient to merely deliver their packets; instead, delivery must also meet a set of class-specific quality of service (QoS) constraints. These include loss rate, delay, and jitter.

IEEE 802.11 [1], the standard for wireless local area networks (WLANs), currently supports nominal data rates up to 54 Mb/s. However, the open shared air medium and use of omnidirectional transmissions place additional burdens on the medium access control (MAC) protocol, compared to wired LANs. 802.11's distributed coordination function (DCF) provides a simple, flexible mechanism for sharing the medium, but lacks the ability to guarantee service levels to multimedia applications.

There has been considerable effort made to improve the MAC's ability to serve and interact with higher-level QoS mechanisms. IEEE's 802.11e Task Group has worked to design and develop a framework for QoS support [2], focusing primarily on providing differentiated access to individual traffic classes (TCs). 802.11e's enhanced DCF (EDCF) uses priority concepts to provide different service levels to each TC. However, EDCF can only provide probabilistic service assurances: TC_A *should* receive better service than TC_B. As service classes are not strictly enforced, the servicing of low-priority traffic can sometimes degrade the service provided to high-priority TCs. Unfortunately, such occurrences are not uncommon, particularly under heavily congested conditions.

These protocols estimate network conditions based on the aggregated network behavior of all traffic classes, reducing the effectiveness of differentiation techniques. Some works have proposed adaptive contention window (CW) schemes, designed to coordinate MAC parameters between different stations. In adaptive EDCF (AEDCF), each node measures the collision rate in order to guide the adjustment of its CW. However, AEDCF still differentiates based on nodes rather than flows. A node's traffic can consist of multiple flows belonging to different TCs, with a large variation in traffic from one node to the next.

In the sliding CW (SCW) scheme, separate CW ranges are maintained for each TC. Therefore, the range from which the backoff counter is selected is dependent on the type of traffic being transmitted. By sliding the CWs, SCW dynamically adjusts to changing conditions, but confines oscillations in throughput and delay. This allows the scheme to provide strict service differentiation and good flow fairness, while still maintaining a high level of channel utilization.

 0890-8044/05/$20.00

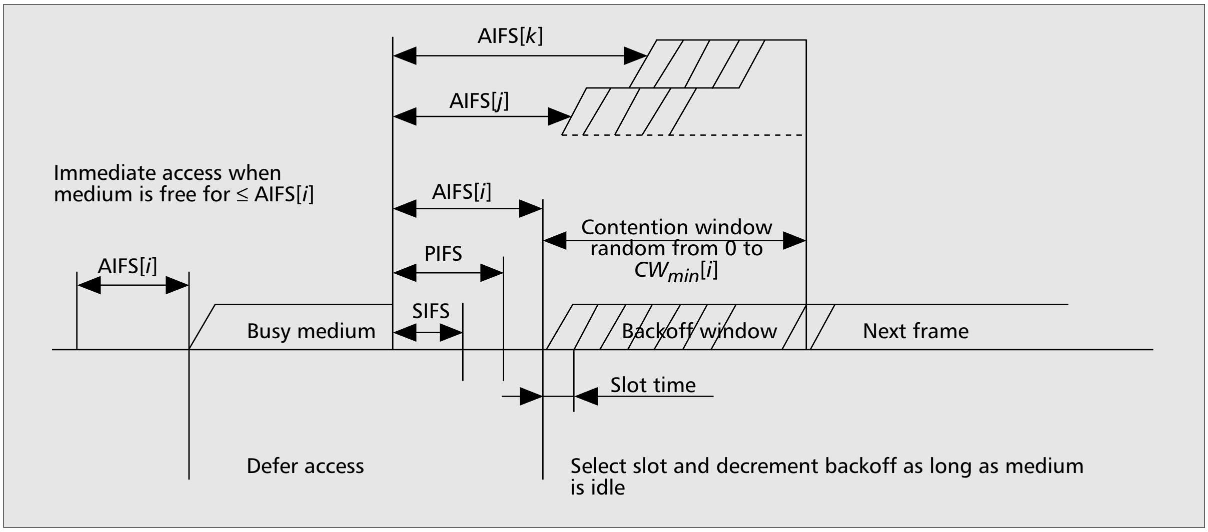

Figure 1. *IEEE 802.11e EDCF channel access.*

The remainder of this article is organized as follows. The next section summarizes related research work on the enhancement of the DCF to better serve real-time traffic. We describe the design of the SCW protocol, and highlight the motivating factors for such an approach. Detailed simulations of SCW have been constructed in ns-2 in order to evaluate its performance under a variety of conditions. These simulations are described, comparing the performance of SCW to both EDCF and AEDCF. Finally, several key conclusions have been drawn from this work; these are stated in the last section.

Background and Related Work on QoS Provisioning for IEEE 802.11

MAC QoS Issues and the IEEE 802.11 DCF

In order to maintain a level of quality acceptable to both the application and the user, packets must be delivered at a sufficiently high rate, in a timely manner. Flows may require different amounts of bandwidth, and may vary in their susceptibility to packet loss. For example, while multimedia applications frequently utilize resilient codecs and application-level error control mechanisms, mechanisms must be in place to ensure that time-sensitive packets are delivered in time. Other applications have a different set of requirements, in terms of bandwidth, packet loss, delay, and jitter.

QoS support frequently utilizes the concept of TCs. Service guarantees can be made to TCs in different ways: absolute guarantees, where the service provided to TC[i] is greater than $c[i]$; or relative guarantees, where the service provided to TC[i] is greater than the service provided to TC[i – 1]. Guarantees, either absolute or relative, can also be made in a probabilistic manner — the service will meet the guarantee with some probability p. The nature of the wireless medium makes absolute guarantees almost impossible to achieve.

IEEE 802.11's DCF [1] is a carrier sense multiple access with collision avoidance (CSMA/CA) scheme, designed to provide contention-based access to the medium. It utilizes an exponential backoff process, doubling the size of the CW after each transmission failure. Backoff intervals are chosen randomly from the range [0–CW]. 802.11's DCF does not provide any means for differentiating TCs. The DCF provides nodes with an opportunity to access the medium, but in fact tends to favor successful transmissions, leading to possible channel domination by a single sender.

Even if QoS mechanisms are added at higher network layers, the MAC must provide sufficient services to support these mechanisms. This typically focuses on either resource reservation or prioritization [3]. However, resource reservation is primarily used with centralized scheduled access, not contention-based protocols.

IEEE 802.11e and the Enhanced DCF

The IEEE 802.11e task group focuses on finding better QoS mechanisms and supporting for multimedia. The 802.11e standard includes an enhanced DCF (EDCF, also known as EDCA), which includes multiple mechanisms for service differentiation. Each node maintains a backoff instance for each TC. MAC service data units (MSDUs) are serviced by the instance for the appropriate TC. This allows each instance to have its own set of contention parameters, specific to that TC. The contention instances for traffic classes i, j, and k are illustrated in Fig. 1.

EDCF includes three mechanisms for the prioritization of traffic classes. The arbitration interframe space (AIFS[i]) replaces the DCF's DIFS), in order to control the time a TC waits before considering the medium to be idle. Minimum and maximum contention window values ($CW_{min}[i]$, $CW_{max}[i]$) allow each TC's window to have different behavior. The persistence factor (PF[i]) provides each TC with its own multiplicative increase factor, compared to the DCF's constant factor of 2. Adjusting these values results in high-priority classes receiving more transmission time than lower-priority TCs [4], without fundamentally changing the conceptual operation of the protocol. However, the use of static parameters creates a scalability problem, as a single parameter set cannot adjust to increasing numbers of active flows. Other works have proposed the use of dynamic parameters in order to allow the MAC to adapt to fluctuating network conditions [4–6]. Adaptive EDCF (AEDCF) [5] extends EDCF by adding a new scheme for resetting the CW size using measurements of the collision rate. This metric guides the adjustment of the CW after each successful transmission. CW changes are smoothed, resulting in slower, less volatile changes. Adaptive fair EDCF (AF-EDCF) [6] takes channel load into consideration by using an adaptive fast backoff decrease mechanism. It also increases the CW while it is deferring whenever it senses the medium becoming busy. This improves protection for high-priority flows.

Adaptive backoff-based differentiation monitors network conditions in order to adjust QoS parameters dynamically [7–9]. However, this is not sufficient for deterministic QoS guarantees, fairness, and bandwidth efficiency. Best effort

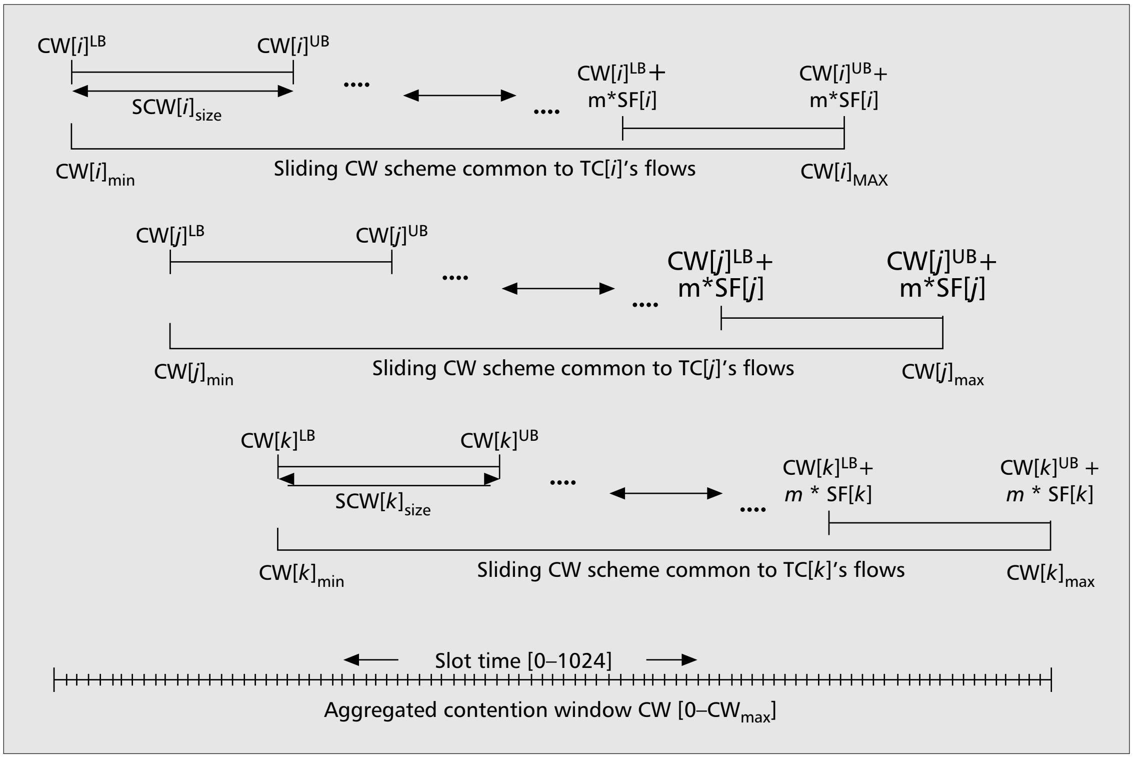

Figure 2. *Sliding contention window scheme for three different TCs.*

traffic may still frequently access the medium, as the actual backoff time is chosen randomly from the interval [0, CW]. This limits the bandwidth available to QoS-sensitive flows. As the number of flows increases and the network becomes congested, priority can be lost due to the random nature of the backoff process.

MAC Fairness

Corresponding TCs should have similar behaviors based on similar MAC parameters. However, two stations can have very different CW values and backoff times for a particular TC, due to collisions and randomness. This can lead to unfairness in medium access opportunities between flows belonging to the same TC. In creating a new MAC scheme to address these QoS issues, inter-TC QoS differentiation must be provided. However, for an effective, efficient, and fair system, intra-TC QoS coordination must also be considered.

Previous work has proposed to monitor the overall network conditions and readjust the CW value of each TC accordingly [5]. However, estimates of metrics such as collision rate can only be made on aggregate network conditions, concerning all TCs. The response to this signal will affect all TCs, which may result in the degradation of perceived QoS, particularly to higher-bit-rate flows. Additionally, the different CW ranges will tend to converge to the same area, resulting in a loss of differentiation and an increase in collisions and delay.

Several approaches [7, references therein] are based on multiplicative-increase linear-decrease (MILD) schemes. These protocols address the fairness problem by including the current CW size within the MAC header of transmitted packets. Nodes overhearing a packet can compare this value with their own, and adjust if necessary. However, as well as increasing the overhead required for the header, this approach suffers from conformance and power issues.

Multimedia service communications (i.e., higher-priority TCs) over WLANs are usually operated and supervised by application-level protocols. For administered WLANs, this likely includes some form of resource allocation, combined with some form of call admission control. Through the remainder of this article, it is assumed that the network has sufficient resources to service all active (accepted) multimedia flows. However, best effort TCs (e.g., Web traffic such as HTTP or FTP) may access the network without any control, creating congestion conditions.

SCW: Sliding Contention Window

Based on previous work, the SCW has been designed: a novel CW scheme that provides QoS differentiation between different classes and fairness between flows of the same class, while still maximizing network utilization. It sustains application-level perceived QoS, guaranteeing the same service levels for all flows within the same TC. By tightening the CW range for each traffic class, SCW eliminates the chance of low-priority traffic receiving a much shorter backoff time than higher-priority traffic. In order to achieve high utilization in low contention periods, the windows are allowed to overlap; however, priority is maintained with only a small amount of random fluctuation. The CWs slide in order to adapt to changes in network conditions in a graceful manner. While this is a considerable change from existing backoff schemes, SCW avoids making further modifications to the 802.11e EDCF.

Service Differentiation Scheme Using Contention Window Ranges

Within the SCW scheme, an SCW[*i*] is associated with each traffic class TC[*i*]. SCW[*i*] has a lower bound $CW[i]^{LB}$ and an upper bound $CW[i]^{UB}$. The lower and upper bounds delimit the interval from which TC[*i*]'s flows select a random backoff value. The bounds of the window change as the window slides, but stay within the interval $[CW[i]_{min}, CW[i]_{max}]$. Figure 2 illustrates the SCW for traffic classes *i*, *j*, and *k*.

All packets from flows belonging to TC[*i*] use the same set of MAC-level parameters, including $CW[i]_{min}$ and $CW[i]_{max}$ as well as AIFS[*i*]. In practice, they can be customized to fit the

Simulation parameters	AIFS (time slots)	SCWsize (time slots)	SF (time slots)	CW_{min} (time slots)	CW_{max} (time slots)	α (maximum tolerated loss)	Packet size (bytes)	Packet generation interval (ms)	Bit rate (b/s)
HP	50	32	16	0	256	5	variable	variable	VBR
MP	70	64	32	32	512	10	1320	40	264000
BE	90	256	128	128	1024	—	812	20	326400

■ Table 1. *Traffic characteristics.*

particular WLAN deployment scenario. Rather than using persistence factors, a TC-dependent sliding factor, SF[i], is used. Whereas PFs represent a multiplicative factor for CW range increases, they create a problem as the lower limit of the window is no longer 0. If PFs were used, the CW would have a dimensioning problem, as its size would change with each window adjustment.

Instead, SCW uses a linear-increase linear-decrease (LILD) model to adjust the SCW range. For instance, each time a flow in TC[i] experiences a high loss rate, SCW[i]'s range is increased by an SF[i] step, until the upper bound reaches the maximum window value. When losses are low and packets are transmitted successfully, rather than resetting the CW, SCW[i]'s range is decreased in the same SF[i] steps until the lower bound reaches $CW[i]_{min}$. The procedures for both decreasing and increasing the SCW range are as follows:

1) ***SCW decreasing procedure***:

IF (oldCW[i]^{LB} — SF[i] ≥ CW[i]$_{min}$) {
 newCW[i]^{LB} = oldCW[i]^{LB} — SF[i]
 newCW[i]^{UB} = oldCW[i]^{UB} — SF[i]
} ELSE {
 newCW[i]^{LB} = CW[i]$_{min}$
 newCW[i]^{UB} = CW[i]$_{min}$ + size(SCW[i])
}

2) ***SCW increasing procedure***:

IF (oldCW[i]^{UB} + SF[i] ≤ CW[i]$_{max}$) {
 newCW[i]^{LB} = oldCW[i]^{LB} + SF[i]
 newCW[i]^{UB} = oldCW[i]^{UB} + SF[i]
} ELSE {
 newCW[i]^{LB} = CW[i]$_{max}$ — size(SCW[i])
 newCW[i]^{UB} = CW[i]$_{max}$
}

At initialization (t_0), each TC i sets its SCW[i] lower and upper bounds according to

$$CW[i]^{LB} = CW[i]_{min}$$
$$CW[i]^{UB} = CW[i]_{min} + 2{*}SF[i]$$
i.e., size(SCW[i])= 2*SF[i].

Note that SCW[i]'s sliding granularity is determined exclusively by the sliding factor for TC[i]. Therefore, SF[i] basically represents the "stride" for adjusting the CW up or down. For our purposes, the higher TC[i]'s priority, the smaller the sliding factor SF[i]. Lower priorities, such as best effort traffic, get larger sliding factors. This ensures that fine-grain adaptation for high-priority traffic occurs, while low-priority traffic is quickly limited when reacting to congested conditions.

Sliding Contention Window Fairness

It is important to remember that constant QoS for each TC is desired, rather than equality between different wireless stations. If traffic were balanced between nodes, achieving fairness between flows within the same TC would require the SCWs on different nodes to remain harmonized. However, as traffic is frequently unbalanced, QoS metric thresholds combined with SCW smoothing rules are used to control the CW's sliding process, in order to ensure that all traffic classes receive the required QoS.

Consider the loss rate $Lr[i]$ of a high-priority flow as perceived by the application. This loss rate accounts for the drop rate measured at the LLC/MAC queue and the frames discarded after successive failed retransmissions. If $Lr[i]$ falls below α_i (a threshold value for the maximum tolerated loss rate for TC[i]), the smoothing rules cause the SCW to be linearly increased in order to give more access opportunity to lower-priority flows. This ensures that the lower-priority TCs still receive adequate QoS, and improves utilization of the medium as well. On the other hand, if the loss rate is too high, the SCW range can be decreased, giving the TC higher priority and reducing the loss rate.

Several LLC/MAC queues are implemented within a single station. Each queue supports one TC, behaving similar to a single DCF entity within the 802.11 standard. As previously mentioned, each queue instance has its own value for AIFS[i], [CW[i]$_{min}$, CW[i]$_{max}$], [CW[i]^{UB}, CW[i]^{LB}], and SF[i]. It is assumed that data from higher layers is tagged with a priority value, by which it is directed into the appropriate queue. Note that a single multimedia stream may be fragmented and mapped onto several different QoS classes.

The queue drop rate ($Lr[i]$) allows us to independently vary the SCW[i] range at every station in order to achieve intra-TC fairness throughout the network. Stations where performance for TC[i] exceeds the required QoS ($Lr[i] \le \alpha_i/2$) increase their CWs, allowing penalized flows of TC[i] (stations having an $Lr[i] \ge \alpha_i$) to gain additional transmission time and consequently reduce their loss rates.

As best effort traffic does not require any QoS metric thresholds, the CW must be adjusted slightly differently. The instantaneous network load $B(T)$ is used to adjust the SCW range. $B(T)$ is the fraction of slots that the medium was observed to be busy out of the previous T slots. This includes all slots where a transmission was successfully completed or a collision occurred. If the network load drops below the threshold $B(T)_{Threshold}$, the SCW range for best effort traffic is decreased. If the load exceeds a throughput saturation threshold $B(T)_{Saturation}$, the SCW range is increased [10]. Based on extensive simulations, appropriate values for $B(T)_{Threshold}$ and $B(T)_{Saturation}$ are 0.7 and 0.9, respectively. The entire sliding algorithm is shown below:

SLIDING ALGORITHM
Sliding for high priority flows i (e.g., EF, AF11, AF12, etc.):
IF ($Lr[i] \ge \alpha_i$) then
 Decrease(SCW[i])
Else IF ($\alpha_i/2 < Lr[i] < \alpha_i$) then
 Maintain current SCW[i]
Else IF ($Lr[i] \le \alpha_i/2$) then
 Increase(SCW[i])

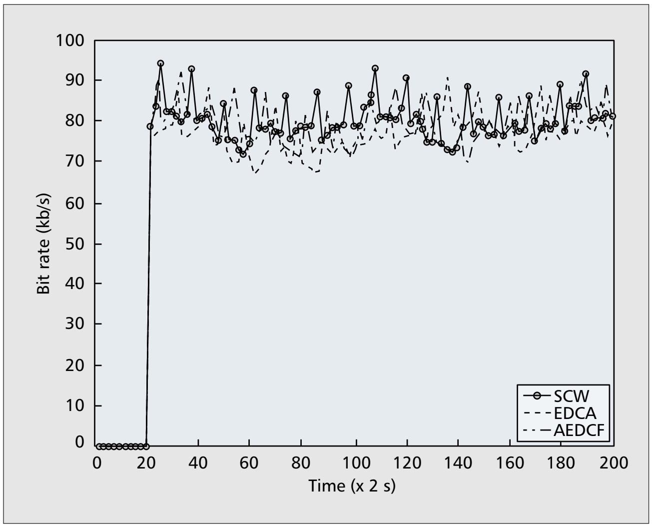

■ Figure 3. *HP flow throughput.*

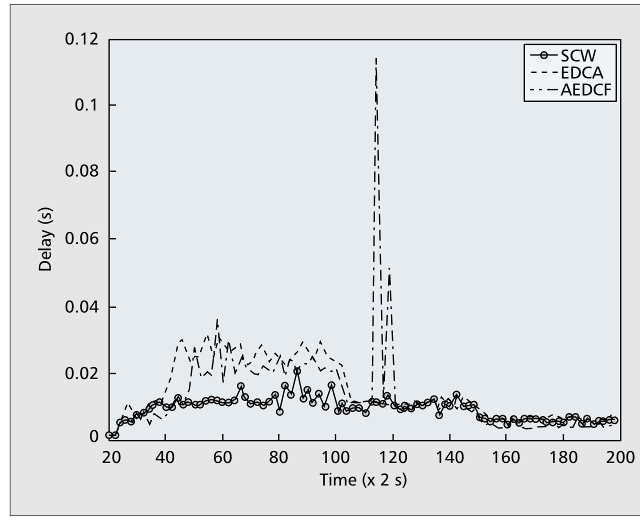

■ Figure 4. *HP flow's delay.*

Sliding for best effort flows k of a given station:
IF ($B(T)/T \le B(T)_{\text{Threshold}}$) then
Decrease(SCW[k])
Else IF ($B(T)_{\text{Threshold}} < B(T)/T < B(T)_{\text{Saturation}}$) then
Maintain current SCW[k]
Else IF ($B(T)/T \ge B(T)_{\text{Saturation}}$) then
Increase(SCW[k])

$B(T)$ is inherently coordinated between stations, as it is calculated based on the overall network load. Each station averages the measurements over the period required to sense CW_{max} idle time slots. By choosing the frequency of measuring $B(T)$ in this way, all backlogged flows are ensured, regardless of priority, to have attempted to access the medium at least once within this period. Therefore, the measurement of $B(T)$ considers all active flows. Throughout this work, the value of $CW[BE]_{max}$ is 1024.

Performance Evaluation

In order to evaluate the advantages of the proposed scheme, we have constructed a simulation of the SCW protocol using ns-2. SCW is compared to EDCF (now called EDCA — enhanced distributed channel access) and A-EDCF using the last IEEE 802.11e draft (IEEE P802.11e/D.8.0, May 2004) [2]. Both SCW and A-EDCF were implemented atop the last NS2 implementation of IEEE 802.11e D8.0 that uses a more realistic MAC implementation with notably a considerably improved backoff freezing process.

The simulations focus on the protocols' abilities to maintain quality of service when many flows are contending for the medium, and to strictly differentiate between service classes. During the simulation, the relative (per-class) network load is deliberately changed to evaluate the ability of SCW (AEDCF and EDCA, respectively) to suit different network configuration.

Simulation Model

For the simulations, the network consists of 10 wireless terminals (WT[i], i = 1, ...,10), and a single access point (AP). WTs are uniformly distributed around the AP, utilizing IEEE 802.11b for communication, at 11 Mb/s. Each WT generates up to *three* different flows at a time, representing three uniquely prioritized traffic classes: high priority (HP), medium priority (MP), and best effort (BE). The properties of these flows are specified in Table 1. During the course of the simulation, the WTs generate a highly dynamic network load in order to investigate SCW's behavior in response to fluctuations in:

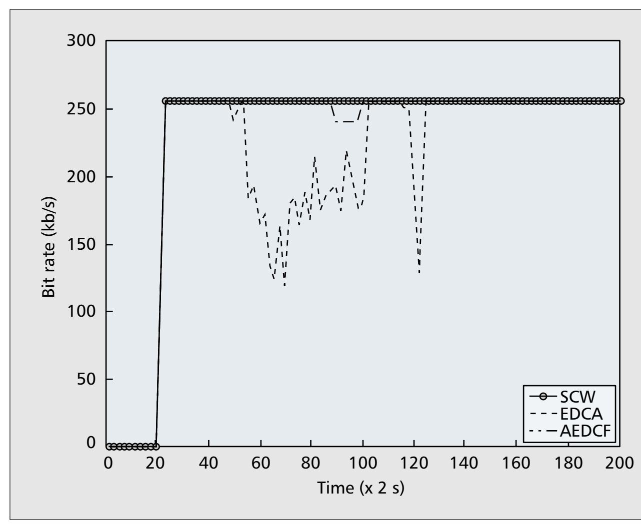

■ Figure 5. *MP flow throughput.*

- The number of active flows
- The overall network load
- The proportional or per-class network load

Each run consists of 200 s of simulated network lifetime. From time t = 0 s to t = 20 s, the channel is empty. Beginning at t = 20 s, new flows of each class are started at 3 s intervals, and begin competing for the channel. By t = 50 s, each WT has three active flows, one of each type, and the wireless channel arrives at saturated conditions, exhibiting a high collision rate. From t = 50 s to t = 100 s, the network remains in this state, to evaluate the degree to which SCW can maintain a strict separation in services. Between t = 100 s and t = 150 s, 5 HP flows and 5 MP flows are stopped. This tests the ability of SCW to maintain high bandwidth efficiency as the network load drops. Finally, 11 additional flows are stopped, consisting of 3 HP, 3 MP, and 5 BE flows. The simulation is completed with 2 HP flows, 2 MP flows, and 5 BE flows remaining active in the network.

Table 1 also shows the MAC-level parameters used in our simulations. EDCA parameters were used as recommended in [2], while AEDCF parameters are those specified in [6]. The updating period for the network measurements (collision ratio) in AEDCF is fixed to 6000 time slots as recommended

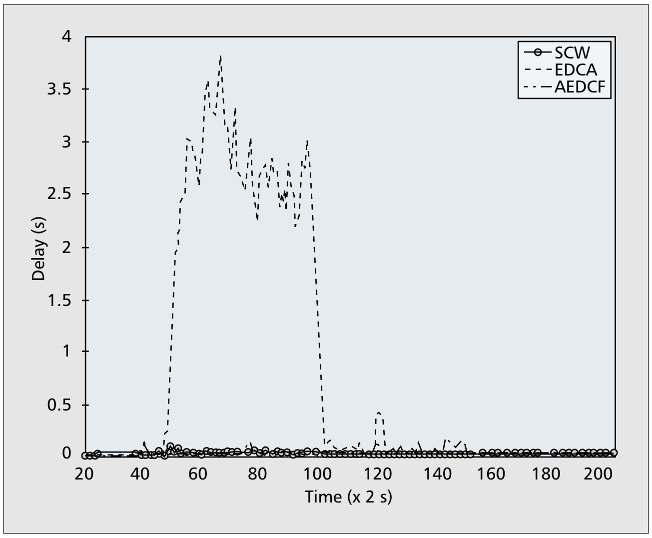

■ Figure 6. *MP flow's delay.*

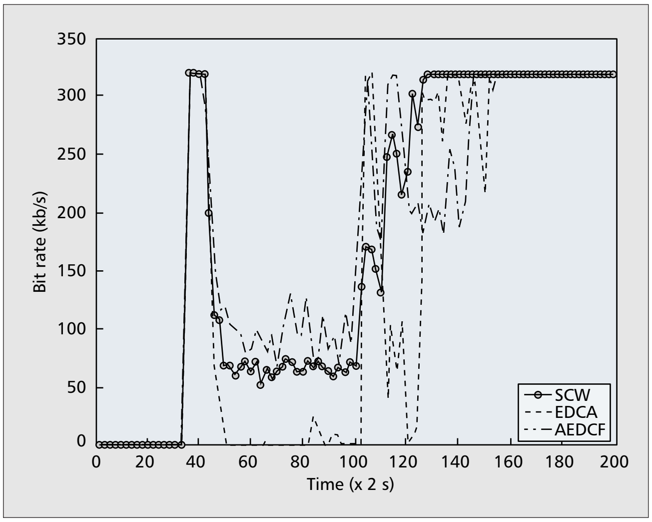

■ Figure 7. *BE flow throughput.*

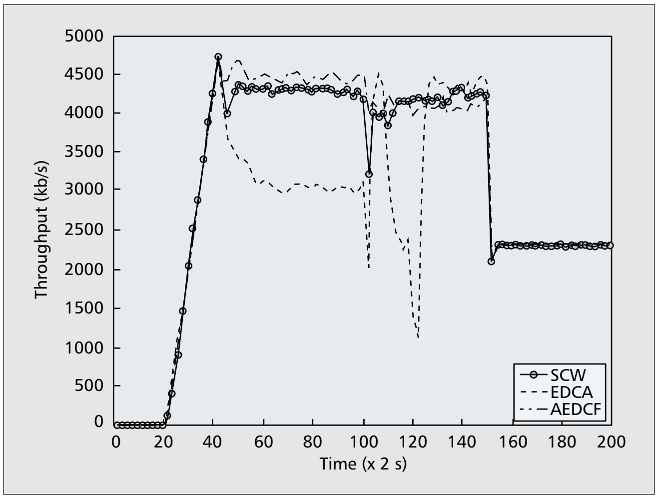

■ Figure 8. *Overall network throughput.*

by the authors. The QoS performance results presented below are measured at the application level, thus representing the perceived QoS. Thus, the measured delays include the queuing delay at the source.

Simulation Results and Analysis

Although a number of traffic scenarios were evaluated, only one can be presented here. In the following, the ability of SCW to provide sustained QoS is assessed for variable bit rate (VBR)-based HP flows using traffic patterns resulting from a real H.264-coded video. Figures 3–8 present simulation results using real H.264 video traces with VBR. The video format is QCIF(176 × 144 resolution), with a mean frame size of approximately 350 bytes. The mean bit rate was approximately 79 kb/s, with a peak rate of 871 kb/s.

As illustrated in Fig. 3, all MAC schemes (SCW, AEDCF, and EDCA) globally succeed in carrying most of the load generated by 10 video traffic sources. Mean throughputs of 72.27, 69.88, and 70.37 kb/s are observed for SCW, AEDCF, and EDCA, respectively. Between $t = 50$ s and $t = 100$ s, SCW achieved a throughput improvement of about 7 kb/s. This is mainly due to high intraclass contention provoked by a too narrow backoff range for HP flows. Actually, as different TCs use different AIFS intervals, the collisions are most likely to occur between flows belonging to the same TC, which entails frequent timeslot waste.

The relatively good performance of both AEDCF and EDCA is due to the use of a scenario where the VBR sources are backlogged with 3 s intervals, which means that the bit rate peaks (871 kb/s) do not occur at the same time. In this particular context, the intra-TC contentions for HP flows are greatly reduced since:

- The peak of each HP flow does not coincide with the peaks of other HP flows.
- The burstiness of each HP flow is amortized by buffering at the MAC level.

In fact, the desynchronization between the throughputs of HP flows within SCW, AEDCF, and EDCA is caused by different queuing times of the offered load. This is more apparent from Fig. 4, where it is clear that in AEDCF and EDCA the HP flows suffer from higher queuing delays at the MAC level.

Figures 5 and 6 give the throughput and delays achieved by MP flows. As expected, both SCW and AEDCF allow MP flows to transmit their offered load throughout the simulation time, although AEDCF's MP flows experience occasional degradations that have consequences for the delay (Fig. 6). Within AEDCF, the MP flows further occupy the channel, increasing the intra-TC collisions and thus the queuing delay at MP flows. This translates into high jitter, with devastating consequences on the perceived video quality at the receiver. This effect is multiplied tenfold with EDCA, where a too narrow backoff range (15–31) provokes high intra-TC contention and a serious drop in throughput.

Within SCW, in addition to reduced throughput oscillations, BE flows achieve a higher mean throughput between $t = 100$ s and $t = 150$ s (Fig. 7). This allows SCW to achieve better overall network utilization during this period. It appears from Fig. 8 that EDCA suffers a devastating drop in network utilization during the period [109–122 s]. As seen from Figs. 5 and 7, MP and BE flows achieve low throughput during this period. Due to the spillover effect of the shared medium, the high intra-TC contention is shifted to affect BE flows. As HP and MP are fairly absorbed by EDCA, excessive transmission (backoff freezing) by BE flows encourages the overflow of MAC queues and leads to BE flows' starvation.

Under stressed conditions, SCW gains a significant advantage over EDCA. The goodput gain of SCW reaches about 40 percent when the load rate is at its maximum (between $t = 100$ s and $t = 150$ s). This represents roughly 1200 kb/s of realized data rate gain.

Between $t = 100$ s and $t = 150$ s, SCW outperforms AEDCF with 100 kb/s in network utilization gain due to higher BE flow throughputs (Fig. 7). Particularly, since BE flows use network load $B(T)$ measurements to adjust (slide up/down) their CW, the channel is more perfectly filled. In contrast, AEDCF uses a per-flow collision rate to adjust the backoff range between 0 and a certain CWx. As illustrated in Fig. 7, this behavior causes excessive throughput oscillations compared to SCW, making it difficult to fill the residual bandwidth. Moreover, each flow may experience widely different collision rates, increasing the disparity between the throughputs achieved by different BE flows.

When using VBR sources the overall network utilization (for both SCW and AEDCF) does not significantly increase during the period [100–150 s]. This is due to low BE throughputs during this period, primarily because BE flows use a too short averaging period for their respective measurements. With VBR sources, those measurements present high short-term fluctuations that depend on:

- The particular burstiness of the sources
- The time when the different bit rate peaks coincide

By increasing the measurement averaging period, higher network utilization may be traded off for poor responsiveness to network configuration changes.

Conclusions and Future Works

This article presents a novel service differentiation scheme based on the separation of contention window ranges. Compared to the EDCF of IEEE 802.11e and AEDCF, the proposed scheme offers an improved ability to differentiate the realized quality-of-service between traffic classes, while remaining fair and bandwidth efficient.

Simulations have shown that the SCW protocol achieves numerous performance gains over the AEDCF and EDCA. In addition to stronger service differentiation between traffic classes, SCW improves throughput and reduces delay and loss rate. It also greatly reduces protocol-induced oscillation effects, and increases fairness between flows of the same traffic class.

QoS-based MAC mechanisms, such as EDCA and AEDCF, must establish an appropriate dimensioning of MAC parameters (e.g., $CW[i]_{min}/CW[i]_{max}$) in respect to the offered load of each traffic class in the network. The chosen parameters may be effectively optimized for a particular network configuration, but perform poorly when the per-class network load changes. This is clearly an issue for a network operator that wishes to offer different service classes and sustain service class guarantees (loss, delay, jitter) regardless of the per-class network load. SCW maintains the differentiation of service classes, despite load variations, without additional configuration by the operator.

Although SCW provides better ability to accommodate the fluctuations in the per-class network load, additional optimizations may be achieved in order to respond to any network. Currently, SCW class-specific MAC parameters must be carefully adjusted, based on the WLAN deployment scenario and predicted traffic. Hence, future work should focus on deriving an analytical model for efficiently dimensioning SCW parameters, as well as finding mechanisms for dynamically (during runtime) adapting these parameters to particular network scenarios.

References

[1] IEEE Std. 802.11, "IEEE, Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications," June 1999.

[2] IEEE 802.11e draft 8.1, "Wireless LAN Medium Access Control (MAC) Enhancements for Quality of Service (QoS)," May 2005.

[3] W. Pattara-Atikom, P. Krishnamurthy, and S. Banerjee, "Distributed Mechanisms for Quality of Service in Wireless LANs," *IEEE Wireless Commun.*, June 2003, vol. 10, no. 3, pp. 26–34.

[4] I. Aad and C. Castelluccia, "Differentiation Mechanisms for IEEE 802.11," *Proc. IEEE INFOCOM '01*, Anchorage, AK, Apr. 2001, vol. 1, pp. 209–18.

[5] L. Romdhani, Q. Ni, and T. Turletti, "Adaptive EDCF: Enhanced Service Differentiation for IEEE 802.11 802.11 Wireless Ad Hoc Networks," *Proc. IEEE WCNC '03*, New Orleans, LA, Mar. 2003, vol. 2, pp. 1373–78.

[6] M. Malli *et al.*, "Adaptive Fair Channel Allocation for QoS Enhancement in IEEE 802.11 Wireless LAN," *Proc. IEEE ICC '04*, Paris, France, July 2004, vol., pp. 347–75.

[7] Z. J. Haas and J. Deng, "On Optimizing the Backoff Interval for Random Access Schemes," *IEEE Trans. Commun.*, vol. 51, no. 12, Dec. 2003, pp. 2081–90.

[8] A. Veres *et al.*, "Supporting Service Differentiation in Wireless Packet Networks using Distributed Control," *IEEE JSAC*, vol. 19, no. 10, Oct. 2001, pp. 2081–93.

[9] F. Cali, M. Conti, and E. Gregori, "Dynamic Tuning of the IEEE 802.11 Protocol to Achieve a Theoretical Throughput Limit," *IEEE/ACM Trans. Net.*, vol. 8, no. 6, Dec. 2000, pp. 785-90.

[10] G. Bianchi, "Performance Analysis of the IEEE 802.11 Distributed Coordination Function," *IEEE JSAC*, Mar. 2000, vol. 18, pp. 535–47.

[11] A. Nafaa, A. Ksentini, and A. Mehaoua "SCW: Sliding Contention Window For Efficient Service Differentiation in IEEE 802.11 Networks," *Proc. IEEE WCNC '05*, New Orleans, LA, Mar. 13, 2005.

[12] A. Nafaa *et al.*, "Towards Backoff Range-Based Service Differentiation Over IEEE 802.11 Wireless LAN Networks," Tech. rep., Univ. of Versailles, Nov. 2004; http://www.prism.uvsq.fr/~anaf/tech_report.pdf

Biographies

ABDELHAMID NAFAA (anaf@prism.uvsq.fr) received a Dipl.-Ing. degree in computer science and information technology from USTHB, Algiers, in 2001. He received his M.Sc. degree in telecommunications and multimedia networking from the University of Versailles, France, in 2002, where he is currently working toward a Ph.D. degree in the area of real-time multimedia communications over mobile and packet-lossy channels. His research interests include wireless IP networking, QoS, and reliable real-time video communications.

ADLEN KSENTINI (kad@prism.uvsq.fr) received a B.S. degree in computer engineering from USTHB, Algiers, in 2000. He received an M.S in telecommunication and multimedia networking from the University of Versailles in 2002. He is currently a Ph.D. candidate at the University of Cergy-Pontoise with the LICP Laboratory. His research interests include QoS schemes for IEEE 802.11 and 802.16 networks, and multimedia over wireless networks.

AHMED MEHAOUA (mea@prism.uvsq.fr) is an associate professor at the University of Versailles, France, and leader of the Multimedia and Mobile Networking Group at the CNRS PRiSM Laboratory. He received his Ph.D. degree in computer science from the University of Versailles in 1997. His primary research interests are multimedia communications and QoS management in wired and wireless broadband networks.

BRENT ISHIBASHI (bkishiba@bbcr.uwaterloo.ca) is currently working towards a Ph.D. at the School of Computer Science of the University of Waterloo. He received his B.Sc. from the University of Guelph (2000), and his M.Math from the University of Waterloo (2004). His research has focused on resource management in multihop wireless networks, particularly mobile ad hoc networks and wireless mesh networks.

YOUSSEF IRAQI (iraqi@bbcr.uwaterloo.ca) is a Research Assistant Professor in the School of Computer Science of the University of Waterloo. He received his B.Sc. from Mohamed V University in Morocco (1995), and his M.S. (2000) and Ph.D. (2003) degrees from the University of Montreal. From 1996 to 1998, he was a research assistant at the Computer Science Research Institute of Montreal. His research interests include network and distributed systems management, resource management in multimedia wired and wireless networks, and peer-to-peer networking.

RAOUF BOUTABA (rboutaba@bbcr.uwaterloo.ca) is an Associate Professor in the School of Computer Science of the University of Waterloo. He has published over 150 refereed papers in the areas of network and distributed systems management, and resource management in multimedia wired and wireless networks. He serves the IEEE Communications Society as a Director on the Board of Governors, and as Vice Chair of the Technical Committee on Information Infrastructure. He also serves on the editorial boards and program committees of numerous journals and conferences.