Quality of Service in Ad Hoc 802.11 Networks

Marcelo G. Rubinstein[1] and José Ferreira de Rezende[2]

[1] Depto. de Eng. Eletrônica e Telecomunicações - Universidade do Estado do Rio de Janeiro
FEN - Rua São Francisco Xavier, 524 - 20550-013 - Rio de Janeiro RJ - Brazil
rubi@uerj.br

[2] Grupo de Teleinformática e Automação - Universidade Federal do Rio de Janeiro
PEE-COPPE - CP 68504 - 21945-970 - Rio de Janeiro RJ - Brazil
rezende@gta.ufrj.br

Abstract

This paper presents major quality of service problems in IEEE 802.11 networks. It analyzes unfairness in bandwidth sharing and TCP instability in MAC 802.11. Moreover, differentiation techniques for providing medium access priority are evaluated. Simulation results show that the ad hoc routing protocol has a great influence on unfairness and on the TCP throughput instability due to succeeding routing failures originated by the MAC sublayer. Differentiation mechanism results show that DIFS-based schemes can provide a fine-grained throughput differentiation compared to backoff-based ones. Nevertheless, backoff-based schemes are more effective in providing latency differentiation, which is important for real-time traffic.

Keywords: Ad hoc networks, IEEE 802.11, Quality of Service

1 Introduction

Wireless networks are being increasingly used in the communication among devices of the most varied types and sizes. Personal computers, handhelds, telephones, appliances, industrial machines, sensors, and others are being used in several environments, such as residences, buildings, cities, forests, and battlefields. Different wireless network standards and technologies have appeared in the last years to accommodate this vast range of applications and coverage.

In this context, the 802.11 [11] standard has obtained an enormous success. For its diversity in terms of capacity and coverage and because of the low cost of network devices, 802.11 has been used in the most varied scenarios. We can find 802.11 in access networks for 2G/3G cellular networks, in residential and campus solutions (local and metropolitan networks), or even in medium distance point-to-point connections in corporate solutions. This vast applicability of the 802.11 has been the key of its commercial success.

The 802.11 standard specifies the Medium Access Control (MAC) protocol and three physical layers for different coverage and speed [13, 12]. Regardless of the continuous capacity increase of these networks, current specifications offer a limited support for Quality of Service (QoS). In ad hoc networks, the distributed access control aims at providing fairness in bandwidth allocation, in which all stations receive the same service independently from their QoS requirements. Protocols and specific mechanisms for 802.11, used in different approaches and architectures, have been proposed to provide QoS in ad hoc networks [18, 2, 10]. By the end of 2000, the 802.11e task group [14] started studying and defining mechanisms to support QoS in the MAC layer, following the approach adopted by the service differentiation architecture of the IETF.

Some inherent problems to wireless networks, such as shared medium, need for hop-by-hop error control, and hidden and exposed terminal problems, hinder QoS provisioning in these networks. This article discusses main problems related to QoS provisioning in 802.11 networks and evaluates QoS related issues, such as fairness and instability. In addition, differentiation mechanisms are compared through simulation.

This paper is organized the following way. Section 2 presents the 802.11 MAC sublayer. Section 3 describes main QoS related problems in the MAC 802.11 and service differentiation techniques for 802.11 networks. Section 4 analyzes unfairness in bandwidth sharing and TCP instability in MAC 802.11 and presents a performance evaluation of service differentiation techniques. Finally, concluding remarks are presented in Section 5.

2 802.11 MAC Sublayer

The MAC sublayer specifies Distributed Coordination Function (DCF) and Point Coordination Function (PCF), as medium access functions [11]. PCF is not addressed in this paper since it cannot be used in ad hoc networks because of its centralized characteristic. DCF is a Carrier Sense Multiple Access with Collision Avoidance (CSMA/CA) and positive acknowledgment mechanism. DCF consists of a basic access method and an optional one that uses Request to Send (RTS) and Clear to Send (CTS) frames.

In the basic access, a station that wants to transmit a frame senses the medium. If the medium is idle for at least a time called Distributed Inter-Frame Space (DIFS), the station transmits. If not, transmission is postponed and a backoff process is initiated. In this backoff phase, a station chooses a random time distributed between zero and the size of the Contention Window (CW) and starts a backoff timer. This timer is periodically decremented after the medium is sensed idle for more than DIFS. The backoff timer is stopped when a transmission is detected. When the timer expires, the station sends its frame.

The receiving station uses Cyclic Redundancy Check (CRC) for error detection. If the frame seems to be correct, the receiver sends an acknowledgment frame (ACK) after sensing the medium idle for a period of time called Short Inter-Frame Space (SIFS). By definition, SIFS is smaller than DIFS. If the transmitting station does not receive the ACK, it schedules a retransmission and enters the backoff process. To reduce the collision probability, the contention window starts with a minimum value given by CW_{min}. After each unsuccessful transmission, the contention window increases to a next power of 2 minus 1, until reaching a maximum predefined value called CW_{max}. CW_{min} and CW_{max} values depend on the kind of physical layer being used. CW_{min} is 31 and CW_{max} is 1023 for DSSS (Direct Sequence Spread Spectrum). Moreover, if a maximum number of retransmissions is reached[1], the frame is dropped. To avoid medium capture, prior to transmitting another frame the the sending station will enter the backoff phase.

The optional access [11] adds to the basic one the use of RTS and CTS frames to avoid problems related to hidden terminals [16]. In this method, carrier sensing can be done by physical and virtual mechanisms. Virtual carrier sensing uses medium reservation by exchanging RTS and CTS frames before transmitting data. RTS and CTS contain information about the destination node and the duration time of the transmission of both data and its ACK frames. The use of RTS and CTS is controlled by an $RTS_{threshold}$. After sensing medium idle for DIFS, the sender issues an RTS to the receiver for medium reservation. The receiver responds with a CTS if it is ready to receive, if the medium is idle for at least SIFS. All stations that listen to RTS, CTS, or both, use the duration information to update the network allocation vector used to virtual carrier sensing. Hidden terminals hear the CTS and defer their transmission to avoid collisions. After receiving CTS and waiting the medium to be idle for SIFS, the sender starts transmitting the frame, as in the basic DCF. If the station does not receive a CTS, it enters the backoff phase and retransmits the RTS.

When using the RTS/CTS handshake, collision may still occur due to the different transmission and interference ranges. A transmission range specifies that transmission is possible within a certain radius around the sender. Within a second radius called interference range, the sender may interfere with other transmission by adding to the background noise [16]. One node may not be able to communicate with another node but transmissions from both nodes may collide at an intermediate node. This because the interference range is about two times the transmission range for WaveLAN wireless systems [21] used in our experiments.

[1] The maximum number of retransmissions for RTS and data frames of length less than or equal to $RTS_{threshold}$ is called ShortRetryLimit and for data frames greater than $RTS_{threshold}$ is known as LargeRetryLimit.

3 802.11 MAC and QoS

This section describes QoS related problems and QoS provisioning issues in 802.11 networks. Unfairness in bandwidth sharing and TCP instability problems are presented. Moreover, service differentiation techniques in 802.11 are described.

3.1 Unfairness

The use of MAC 802.11 in high load conditions may cause unfairness problems. DCF must provide a fair bandwidth sharing, but unfairness in medium access, mainly in multi-hop networks, is reported in [17, 19, 18, 5, 22].

Gerla et al. [17, 20] demonstrate that the interaction between TCP and MAC backoff timers in multi-hop networks causes unfairness due to medium capture, i.e., some stations get access to the medium while others not. When two stations contend for the medium and one is "pushed back" by timeouts, the

binary exponential backoff nature of both MAC and TCP timeouts make the situation progressively worse for the looser [17], i.e., the binary exponential backoff always favors the last succeeding station [3]. Bensaou et al. [3, 19], Vaidya et al. [18], and Agrawal et al. [5] also describe this kind of unfairness problem. Agrawal et al. [5] present a mechanism that uses disjoint forward and reverse paths for TCP data and ACK packets and a dynamic contention-balancing scheme to improve TCP performance by avoiding capture conditions. Goff et al. [7] describe a preemptive scheme for on-demand routing protocols that improves TCP performance. This mechanism initiates a new path discovery before a link is broken, hence avoiding TCP backoff timer expiration.

Xu and Saadawi [21, 22] present a TCP problem in multi-hop 802.11 networks called neighboring node one-hop unfairness. This problem is due to one node failing to reach its adjacent nodes. In this case, this node will drop all queued packets and will report a route failure. Before a route is found, no packet can be sent out, what will probably cause a timeout in the TCP sender. Routing failure is rooted in the MAC sublayer because of collisions concerning routing messages and RTS frames, as we will see in Section 4.1.

3.2 Instability

Using TCP over multi-hop wireless networks may cause instability, i.e., throughput may vary significantly even if a single connection is set up and network condition does not vary [22]. By analyzing simulation traces, Xu and Saadawi [21, 22] found that this problem is due to one node failing to reach its adjacent node. One solution to this problem is to limit packet transmission without acknowledgment, i.e., by reducing maximum window size of the transmitter [22].

3.3 Service Differentiation in the MAC 802.11

A number of recent proposals allow service differentiation among stations or even among traffic classes, in the 802.11 standard. This differentiation is achieved by assigning different priorities in the wireless medium access to stations that contend for it. These proposals suggest modifications to the DCF mode.

These techniques can be classified according to the parameter used to achieve differentiation: DIFS, backoff, frame size, and RTS/CTS threshold.

The DIFS-based scheme consists of configuring wireless stations with different values for this parameter according to the priority that one wishes to assign to each station. The larger the DIFS, in the number of slots, the smaller the station priority. To avoid contention among stations with different priorities, the maximum contention window of a station with priority j added to DIFS_j is chosen in such a way that it is never larger than DIFS_{j+1} (lower priority). This guarantees that a higher priority station has no frames to send when a lower priority station starts transmitting [1].

The backoff-based scheme consists of assigning different intervals (min and max) for the contention window of each station or determining how the contention window evolves along with station/flow priority, number of retransmission retrials, and other factors. In [18], the calculated backoff interval is proportional to the frame size to be sent and inversely proportional to the weight of the flow to which the packet belongs. Therefore, stations with greater weights (higher priority) choose smaller backoff times. The inclusion of the frame size provides a weighted fair share of the bandwidth. Thus, this scheme implements, in a distributed way, the weighted fair-queuing scheduling mechanism. In [1], the contention window intervals are calculated according to the priority established for each station. Aad et al. [2] also present a mechanism that assigns different priorities for different destinations, i.e., per-flow differentiation. In [9, 10], the authors propose a scheme where the priority of the next frame to be sent is included in RTS and CTS control frames, data frame, and the corresponding ACK. Since all stations in the same coverage area hear this information, they can maintain a table with the current head-of-line frames of all stations that contend for the medium. The contention window interval is calculated then by each station according to the position (rank), in terms of priority, of its frame in that table. Bensaou et al. [19, 3] propose a scheme of differentiated backoff in which each station adjusts its contention window according to the estimate of its bandwidth share and the share obtained by the other stations. The main idea is to allow all stations to transmit using the default configuration if the total load is smaller than the link capacity. In case of exceeding the link capacity, each station should obtain an access proportional to a sharing index previously established in the admission control. In the next section, we evaluate the degree of differentiation one can obtain by using the two schemes described before.

The two schemes described below establish a coarser differentiation. In the technique based on the frame size, stations with higher priority use larger frame sizes in their transmissions. This scheme con-

trols the time a station retains the medium after winning a contention for it.

The technique based on the RTS/CTS threshold consists of the use of medium reservation through the RTS/CTS handshake. Stations with thresholds values larger than frame sizes of a certain flow will not use RTS/CTS. These frames will have higher collision probability and consequently a lower priority.

The 802.11e [14, 8] introduces a new coordination function, called HCF (Hybrid Coordination Function), which provides contention-based and contention-free access modes. In the contention-based access, called Enhanced DCF (EDCF), the contending flows use different DIFS and CW_{min} settings.

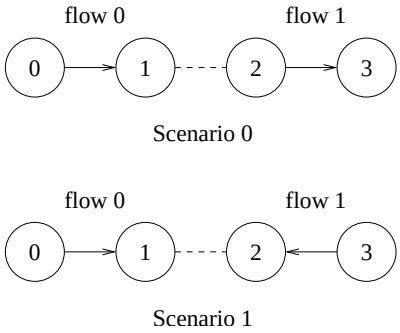

Figure 1: Scenarios used in simulations.

4 Simulation Results

In this section, we aim at providing an exhaustive analysis of the unfairness and instability problems discussed above. Moreover, we are interested in evaluating the differentiation level offered by modifying MAC parameters, such as DIFS and contention window, in different scenarios. The Network Simulator ns-2 [6] is used in the simulation studies. We have used the functionalities of 802.11 networks added with service differentiation, ad hoc routing protocols, and CBR and TCP traffic sources.

4.1 Unfairness Problem

The first analysis is related to the unfairness problem. Figure 1 presents two scenarios used in this experiment. Arrow lines indicate traffic between stations and dashed lines mean that stations are in the communication range of each other but there is no traffic between them. We use static nodes and a small topology because we need to limit link failures in order to be able to present an in-depth analysis of the unfairness problem. Nodes are 200 m far from each other, what allows a node to directly communicate only with its neighbor nodes. The channel capacity is 11 Mbps. The nominal transmission range is about 250 m and the nominal interference range is about 500 m.

As pointed out by Xu and Saadawi [22], TCP traffic increases the problems in the MAC sublayer. In order to evaluate the unfairness problem without considering this influence, CBR sources are used. Both CBR flows generate 1000-byte packets and the interval between these packets is varied in the simulations. Communication between nodes 0 and 1 (flow 0) starts at 0 s and flow 1, between nodes 2 and 3, starts at 5 s. Figures 2(a) and 2(b) show the throughput obtained by these two flows within each 5 s time interval for both DSDV (Destination-Sequenced Distance Vector) [15] and DSR (Dynamic Source Routing) [4] routing protocols, respectively.

For DSDV and a CBR rate of 1500 kbps (Figure 2(a)), several packets are dropped at the beginning of the communication because node 0 is short of buffer space on the interface queue between MAC and logical link sublayer. The reason for that is the lack of route from node 0 to node 1. After obtaining all routes, some DSDV routing packets, which are sent in broadcast with no RTS/CTS handshake, from node 3 and RTS frames from node 0 to node 1 are dropped because of collisions between them, i.e., broadcast transmissions from node 3 interfere with unicast transmissions from node 0 to node 1 and vice versa. In addition, ACK frames from node 1 to node 0 cause dropping of DSDV routing packets on node 2. These dropped routing packets do not affect the throughput of flow 1 since in the ns-2 implementation [4] the number of periodic update missed before a link is declared broken[2] is 3, which is never reached for this experiment. Figure 2(b) shows the throughput for DSR. For this protocol, link breakage detection feedback from the MAC sublayer is used. The MAC sublayer signals when the maximum number of retransmissions is reached. In this experiment, few RTS frames from nodes 0 and 2 are dropped because of collisions. As a consequence there is no link breakage. Therefore, when the aggregate rate of the CBR flows is not higher than the maximum achievable throughput, which is around 3.6 Mbps for this experiment, fairness is achieved in the bandwidth allocation.

For a CBR rate of 3000 kbps, several data packets are dropped at the interface queue, since the aggregate

[2] In the ns-2 implementation, for DSDV there is no link breakage detection from the 802.11 MAC protocol [4].

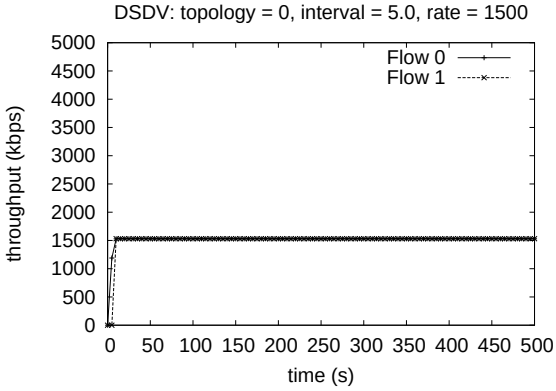

(a)

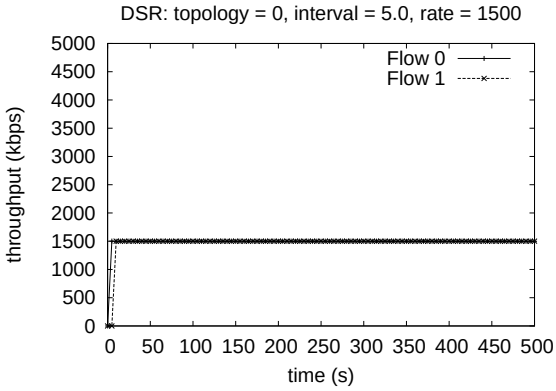

(b)

Figure 2: Throughput for scenario 0 and CBR rate of 1500 kbps.

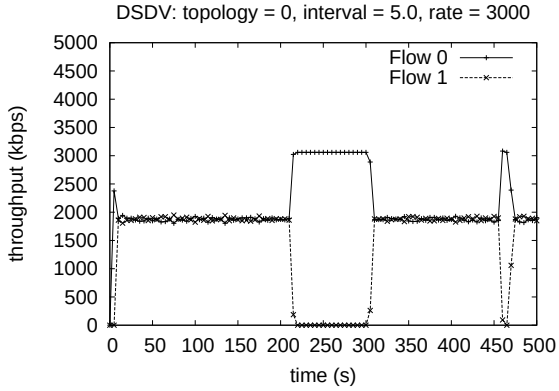

(a)

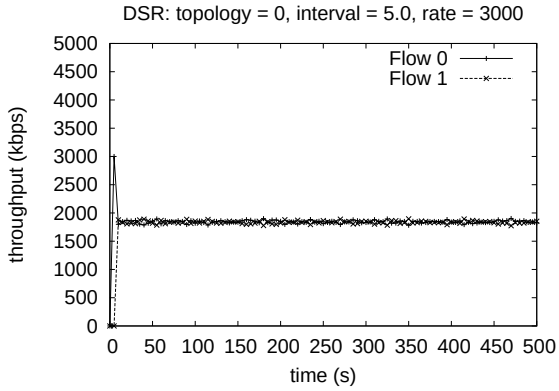

(b)

Figure 3: Throughput for scenario 0 and CBR rate of 3000 kbps.

rate is higher than the maximum achievable throughput. When DSDV protocol is used (Figure 3(a)), channel sharing is not fair. For example, from 210 s to 305 s, DSDV routing messages from node 3 are dropped on node 2 because of collisions and a link breakage is declared. Since node 2 has no route to send packets to node 3 and the maximum number of packets buffered per node per destination is five in ns-2 [4], several packets are dropped at the interface queue until a new route is found. This occurs again at 455 s. Figure 3(b) shows a fair bandwidth sharing for DSR.

The main reason for unfairness at 4500 kbps rate when DSDV is used (Figure 4(a)) is the number of routing messages which are dropped on node 2 because of collisions with node 0 transmissions (data or RTS frames). Unfairness is much greater because the aggregate rate is higher. For DSR (Figure 4(b)) besides dropped packets at the interface queue, there are also packets from node 2 that expire after waiting for a route from 5 s to 25 s.

We observe in the scenario 0 that flow 0 is always favored against flow 1. This because large data frames from flow 0 compete with small ACK frames from flow 1. Therefore, scenario 1 has been used to assess the throughput performance when data frames from both flows may collide. In this scenario, retransmissions of RTS frames are much more frequent than for

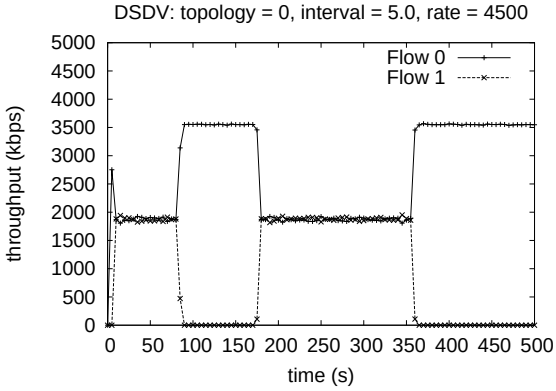

(a)

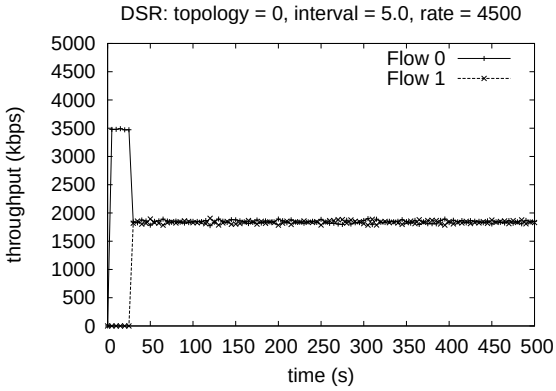

(b)

Figure 4: Throughput for scenario 0 and CBR rate of 4500 kbps.

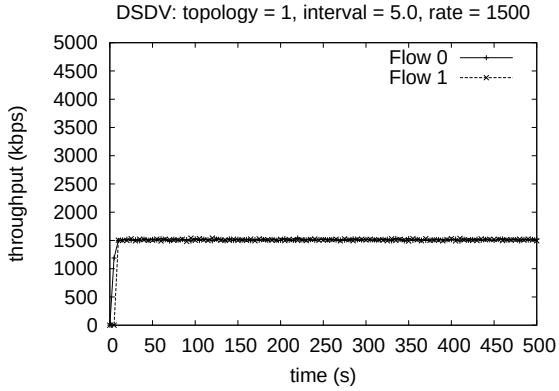

(a)

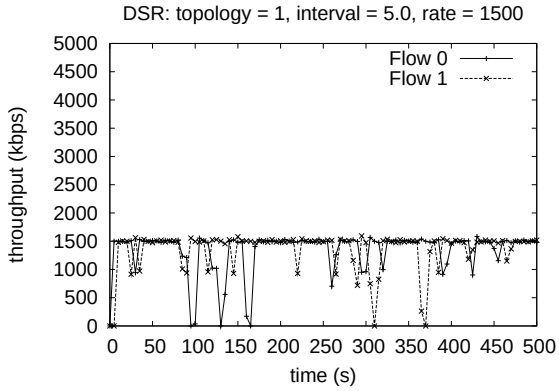

(b)

Figure 5: Throughput for scenario 1 and CBR rate of 1500 kbps.

scenario 0 because data are sent from node 3 to node 2 and neither node 0 nor node 1 can listen to node 3 data frames. The same problem happens for node 0 transmissions. We expect performance reductions for both routing protocols.

For DSDV (Figure 5(a)), channel sharing is fair but many data packets are dropped at nodes 3 and 0. For example, when node 3 (or node 0) fails ShortRetryLimit times, which is seven in ns-2, to reach node 2 (or node 1), the packet is dropped. For DSR, when the maximum number of retransmissions is reached, a link breakage occurs. For example, in Figure 5(b) at 302 s, the route from node 3 to node 2 is lost. As a consequence, many node 3 packets expire before a new route is discovered. Route is reestablished at 312 s. The same happens to flow 0, for example at 90 s.

When the aggregate rate is higher than the maximum achievable throughput for DSDV (Figures 6(a) and 7(a)), packets are dropped, most of them at the interface queue. Frequently one station captures the channel when the other fails to reach its receiver. For DSR (Figures 6(b) and 7(b)), when the CBR rate increases, unfairness grows due to the great number of expired packets.

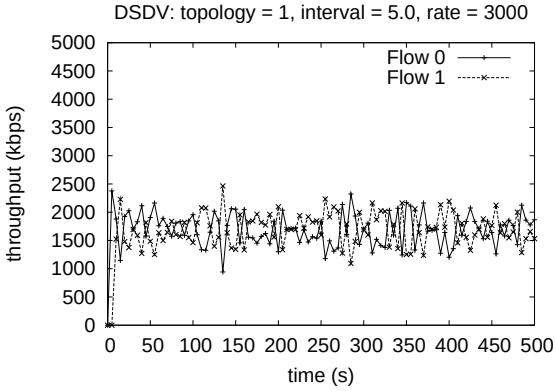

(a)

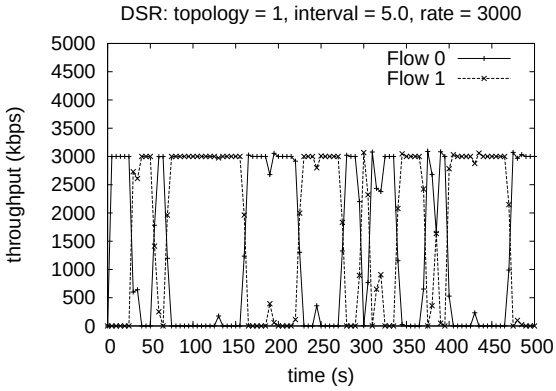

(b)

Figure 6: Throughput for scenario 1 and CBR rate of 3000 kbps.

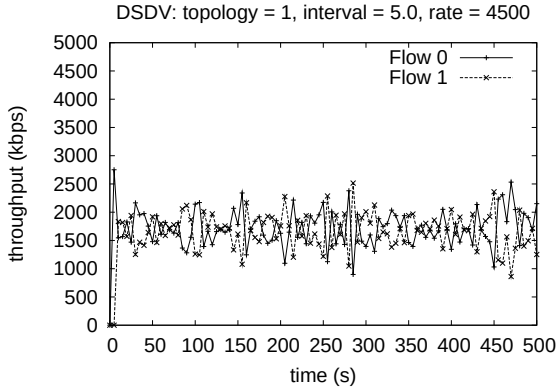

(a)

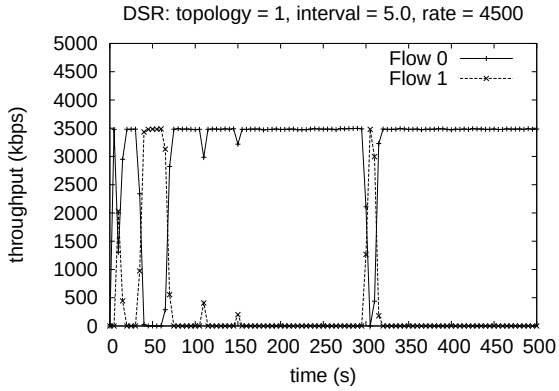

(b)

Figure 7: Throughput for scenario 1 and CBR rate of 4500 kbps.

4.2 Instability Problem

Simulations related to the TCP instability problem have also been performed on the four-node topology used in the previous experiments. A TCP connection is set up between nodes 0 and 3. The TCP packet size is 1 kbytes and an FTP application, which always has data to send, is used. The channel capacity is 2 Mbps and the DSR protocol is used. TCP maximum transmission window size is varied in the simulations. Throughput is plotted over 1.0 s intervals.

As the network condition does not vary, TCP throughput should stay stable around an operation point. Figure 8(a) shows that for a window size of 64, there are intervals on which node 0 does not receive any ACK. For a window size of 4 (Figure 8(b)), node 0 always receives ACKs. According to simulation traces, there are RTS collisions due to the different transmission and interference ranges. Most of these collisions occurs at node 2 when node 3 sends an RTS to node 2 (node 3 wants to send a TCP ACK) and node 0 is sending a TCP data frame or an RTS to node 1. Since for this experiment the nominal transmission range is about 250 m and the nominal interference range is about 500 m, node 3 does not listen to the transmission of node 0 and collisions occur (hidden terminal problem). When node 3 fails to reach node 2 ShortRetryLimit times, the MAC sublayer reports a link

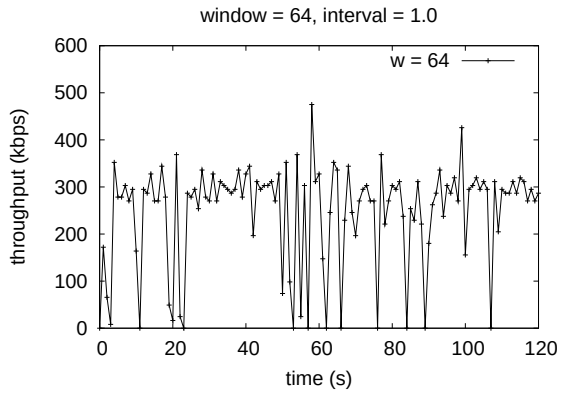

(a)

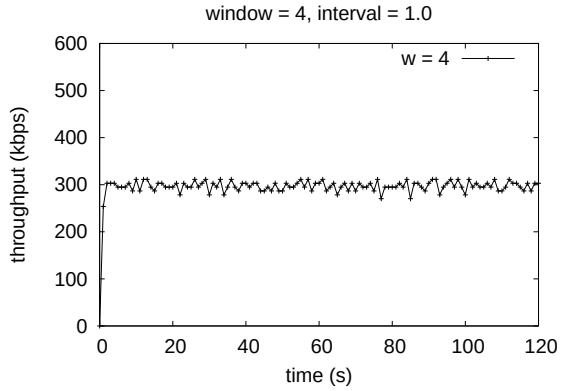

(b)

Figure 8: Throughput with TCP window size of 64 (a) and 4 (b).

breakage to its upper layer and a route failure event is notified. Before a route is found, no data packet can be sent out and usually this causes a timeout in the TCP sender, degrading TCP performance. TCP in Figure 8(b) does not present this instability problem because the maximum number of packets sent at one time (back-to-back) is four. There is a lower probability of an intermediate node failing to get the medium. Hence, a link breakage due to this failure is less likely.

4.3 Differentiation

The topology used in the differentiation experiments consists of three 802.11 stations transmitting to a fourth station (Figure 9). All stations generate a 1.8 Mbps CBR traffic with packet sizes of 1 kbytes. The distance between transmitting stations and the receiving station is 50 m, i.e., within the transmission range of 250 m. The ad hoc routing protocol is DSDV and the channel capacity is 11 Mbps. The maximum achievable throughput in this channel is largely dependent on the frame size used by the sources, the use of the RTS/CTS handshake, the number of stations contending for the medium, and differentiation parameters such as DIFS and the contention window interval. For instance, when only one CBR source contends for the medium and uses packets of 1 kbytes and default values for differentiation parameters, the maximum achievable throughput is 3.6 Mbps. Stations 1, 2, and 3 start their transmissions at 0 s, 10 s, and 20 s. The throughput obtained by each station at each 1 s interval is evaluated.

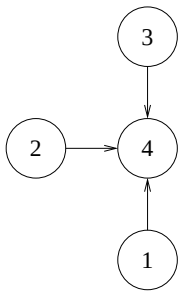

Figure 9: Scenario used in differentiation simulations.

In the first scenario, no differentiation mechanism is used, i.e., all stations have the same default configuration for differentiation parameters. Figure 10(a) shows the throughput obtained by each station. The channel capacity is fairly shared among all stations.

Following, the differentiation level provided by the DIFS-based scheme is analyzed. The default value for DIFS in ns-2 is 2 slots plus SIFS. Stations 1, 2, and 3 have their DIFS configured with 2, 4, and 6 slots plus SIFS, respectively. Figure 10(b) shows the throughput differentiation. Between 10 s and 20 s, only two stations fairly share the channel because their aggregate rate is inferior to the maximum achievable throughput. When the third station starts transmitting, the channel capacity is lower than the total traffic and differentiation starts. Station 1 obtains more bandwidth than stations 2 and 3 because it has the smallest DIFS.

We also evaluate the performance gain of changing the contention window size. Transmitting stations have their contention window intervals

([CW_{min},CW_{max}]) configured to [31:1023], [63:2047], and [127:4095]. DIFS is the same for all stations. In Figure 10(c), the station with the smallest minimum contention window (CW_{min}) obtains more bandwidth.

We also evaluate the degree of throughput control one can obtain by modifying the differentiation parameters. A CBR flow and a long-term TCP connection from station 1 to station 4 are used. Both connections use 576-byte packets. Three different experiments have been performed. One provides no differentiation and the others modify DIFS and minimum contention window size, respectively. DIFS and the contention window interval are kept constant when no differentiation scheme is used. For DIFS, the number of slots is increased by eight slots every second. CW_{min} assumes the value of the next power of two minus one slot every 5 s.

Figures 11(a) and 11(b) show the throughput variation. We observe a linear reduction of the goodput achieved by both connections under the DIFS-based scheme. We observe a drastic reduction of the throughput every time the CW_{min} is increased.

In order to evaluate the latency differentiation in the medium access, Table 1 shows the 90th and 95th-percentile values of latency for each flow in the three cases. Results show that the contention window differentiation scheme is more effective in providing latency differentiation.

Table 1: Latency differentiation.

scheme	station	90th perc. (ms)	95th perc. (ms)
no diff.	STA1	371.27	393.15
	STA2	371.66	388.92
	STA3	384.83	405.05
DIFS	STA1	295.26	307.15
	STA2	381.43	396.56
	STA3	496.19	515.74
cw	STA1	7.91	9.67
	STA2	339.04	346.01
	STA3	783.52	806.76

4.3.1 Increasing the number of stations

In the following scenario, we evaluate the level of service differentiation achieved when the number of stations that share the same wireless channel is increased. All stations transmit to the same receiving station. Three different service classes are used. In the DIFS-based scheme, stations in the service class 1, 2 and 3 use DIFS equal to 2, 4, and 6 slots plus

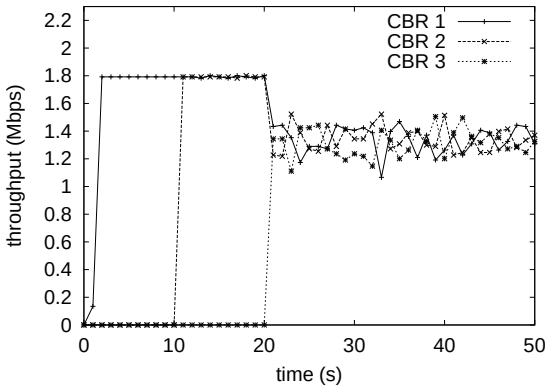

(a) No differentiation

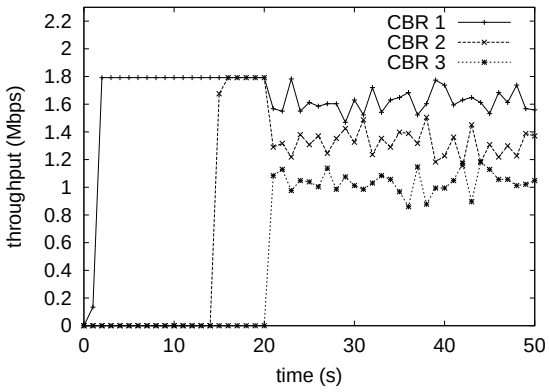

(b) DIFS

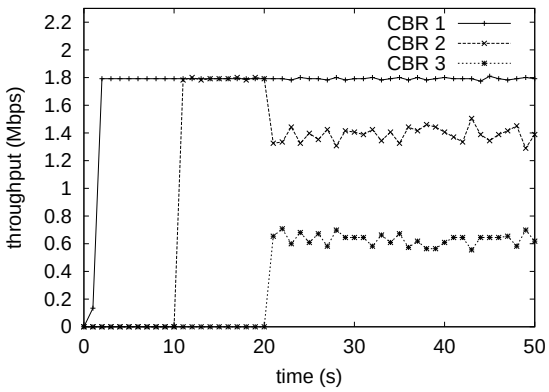

(c) CW

Figure 10: Throughput differentiation.

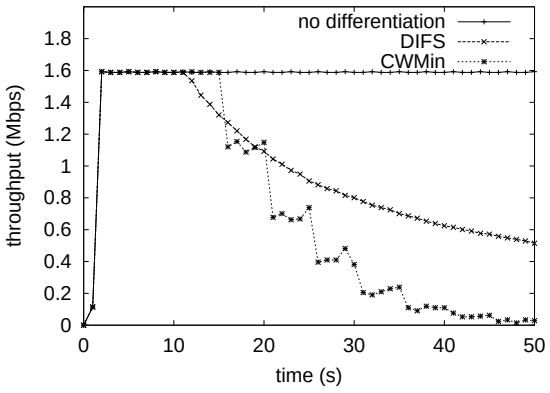

(a) CBR

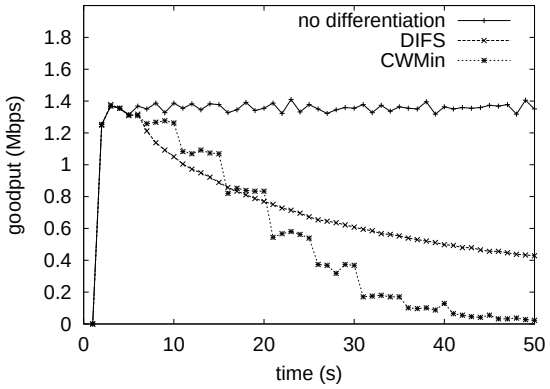

(b) FTP

Figure 11: Throughput control by varying DIFS and contention window intervals.

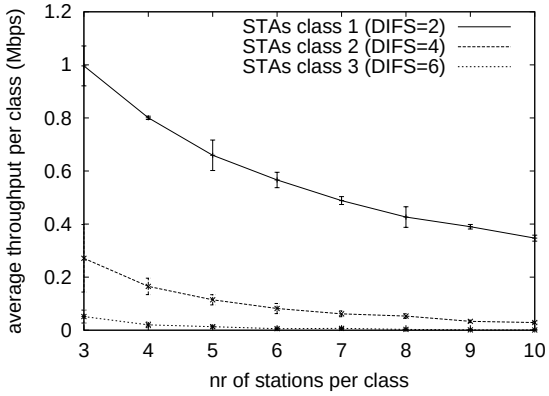

(a) DIFS

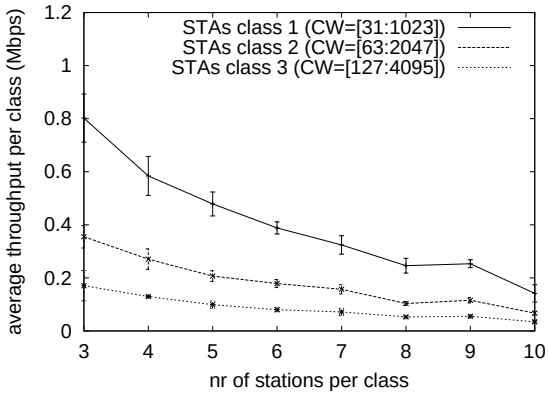

(b) CW

Figure 12: Average throughput when increasing the number of stations per service class.

SIFS, respectively. In the contention window scheme, transmitting stations in the service class 1, 2 and 3 have their contention window intervals configured to [31:1023], [63:2047], and [127:4095], respectively.

Figures 12(a) and 12(b) show the average throughput of stations that belong to one of the service classes. The number of stations per class is increased from 3 to 10. Figures also show 95% confidence intervals. The same throughput differentiation is achieved with the presence of an increasing number of stations.

4.3.2 Adding mobility to the stations

We now evaluate the influence of mobility in the differentiation level achieved by the different schemes. In this scenario, all stations follow a random generated movement pattern, called random-way point [6], in which transmitting stations always stay within the range of the receiving station. Average speed of each station is 10 m/s with movement pauses of 20 seconds in average. The number of stations is increased in each simulation run as described in Section 4.3.1.

Figures 13(a) and 13(b) show the average throughput obtained by stations within the same service class, as described before. Results show that the through-

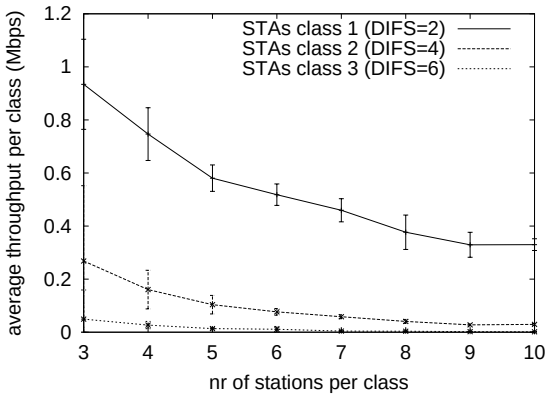

(a) DIFS

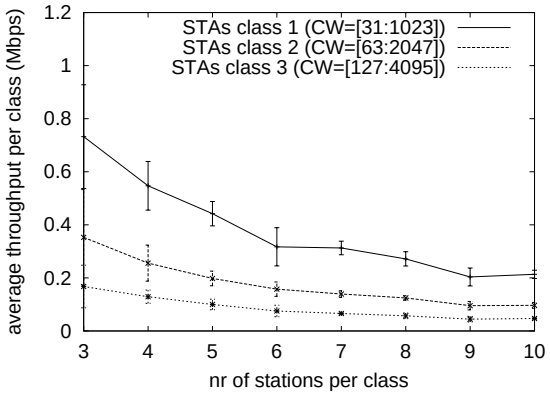

(b) CW

Figure 13: Average throughput for mobile stations.

put differentiation takes effect even in the presence of mobility, however, larger confidence intervals are observed. The average speed and pause have been varied in different simulation runs. The same differentiation behavior is obtained.

5 Conclusion

We have presented major problems in providing QoS in IEEE 802.11 networks. Unfairness in the allocation of channel capacity, TCP throughput instability, and service differentiation have been analyzed by simulation.

Ad hoc network scenarios have been created to assess the fairness in the allocation of bandwidth between different stations and the TCP performance. Results show that ad hoc routing protocols have a great influence on the unfairness and on TCP throughput instability due to succeeding routing failures originated by the MAC sublayer.

DIFS and backoff-based schemes for service differentiation have been evaluated. These approaches have been applied to provide a station-based differentiation instead of a per-flow differentiation. We have assessed the influence of the number of stations and of mobility on throughput and latency differentiation. Results show that DIFS-based schemes can provide a fine-grained throughput differentiation when compared to backoff-based ones. Backoff-based schemes, however, are more effective in providing latency differentiation, which is important for real-time traffic.

Acknowledgments

This work has been supported by UFRJ, FAPERJ, CNPq, CAPES, and COFECUB.

References

[1] I. Aad and C. Castelluccia. Differentiation mechanisms for IEEE 802.11. In *Infocom'2001*, Anchorage, Alaska, April 2001.

[2] I. Aad and C. Castelluccia. Remarks on per-flow differentiation in IEEE 802.11. In *Proceedings of the European Wireless*, Florence, Italy, February 2002.

[3] Brahim Bensaou, Yu Wang, and Chi Chung Ko. Fair medium access in 802.11 based wireless ad-hoc networks. In *IEEE/ACM The First Annual Workshop on Mobile Ad Hoc Networking & Computing (MobiHoc)*, Boston, USA, August 2000.

[4] Josh Broch, David A. Maltz, David B. Johnson, Yih-Chun Hu, and Jorjeta Jetcheva. A performance comparison of multi-hop wireless ad hoc network routing protocols. In *ACM The Forth Annual Conference on Mobile Computing and Networking (MobiCom)*, Dallas, USA, October 1998.

[5] Carlos de M. Cordeiro, Samir R. Das, and Dharma P. Agrawal. COPAS: Dynamic contention-balancing to enhance the performance of TCP over multi-hop wireless networks. In

The 11th International Conference on Computer Communications and Networks (IC3N), Miami, USA, October 2002.

[6] Kevin Fall and Kannan Varadhan. The NS Manual. Technical report, The VINT Project, January 2002.

[7] Tom Goff, Nael Abu-Ghazaleh, Dhananjay Phatak, and Ridvan Kahvecioglui. Preemptive routing in ad hoc networks. *Journal of Parallel and Distributed Computing*, 63(2):123–140, February 2003.

[8] Daqing Gu and Jinyun Zhang. QoS enhancement in IEEE 802.11 wireless local area networks. *IEEE Communications Magazine*, 41(6):120–124, June 2003.

[9] V. Kanodia, C. Li, A. Sabharwal, B. Sadeghi, and E. Knightly. Distributed multi-hop scheduling and medium access with delay and throughput constraints. In *ACM The Seventh Annual Conference on Mobile Computing and Networking (MobiCom)*, Rome, Italy, July 2001.

[10] Vikram Kanodia, Chengzhi Li, Ashutosh Sabharwal, Bahareh Sadeghi, and Edward Knightly. Distributed priority scheduling and medium access in ad hoc networks. *Wireless Networks*, 8(5):455–466, September 2002.

[11] Institute of Electrical and Eletronics Engineers. Wireless LAN medium access control (MAC) and physical layer (PHY) specifications. IEEE Standard 802.11, 1999.

[12] Institute of Electrical and Eletronics Engineers. Wireless LAN medium access control (MAC) and physical layer (PHY) specifications: High-speed physical layer in the 5 GHz band. IEEE Standard 802.11a, 1999.

[13] Institute of Electrical and Eletronics Engineers. Wireless LAN medium access control (MAC) and physical layer (PHY) specifications: Higher-speed physical layer extension in the 2.4 GHz band. IEEE Standard 802.11b, 1999.

[14] Institute of Electrical and Eletronics Engineers. Medium access control (MAC) enhancements for quality of service (QoS). IEEE Draft for Standard 802.11e/D3.2, March 2002.

[15] Charles Perkins and Pravin Bhagwat. Highly dynamic destination-sequenced distance-vector routing (DSDV) for mobile computers. In *ACM SIGCOMM'94 Conference on Communications Architectures, Protocols and Applications*, pages 234–244, 1994.

[16] J. Schiller. *Mobile Communications*. Addison-Wesley, 1^{st} edition, 2000. ISBN 0201398362.

[17] Ken Tang and Mario Gerla. Fair sharing of MAC under TCP in wireless ad hoc networks. In *IEEE Multiaccess, Mobility and Teletraffic for Personal Communications (MMT'99)*, Venice, Italy, October 1999.

[18] Nitin H. Vaidya, Paramvir Bahl, and Seema Gupta. Distributed fair scheduling in a wireless LAN. In *ACM The Sixth Annual Conference on Mobile Computing and Networking (MobiCom)*, Boston, EUA, August 2000.

[19] Yu Wang and Brahim Bensaou. Achieving fairness in IEEE 802.11 DFWMAC with variable packet lengths. In *IEEE Global Telecommunications Conference (Globecom)*, Texas, USA, November 2001.

[20] Kaixin Xu, Sang Bae, Sungwook Lee, and Mario Gerla. TCP behavior across multihop wireless networks and the wired internet. In *Proceedings of the 5th ACM International Workshop on Wireless Mobile Multimedia*, pages 41–48, Georgia, USA, September 2002.

[21] Shugong Xu and Tarek Saadawi. Does the IEEE 802.11 MAC protocol work well in multihop wireless ad hoc networks? *IEEE Communications Magazine*, 39(6):130–137, June 2001.

[22] Shugong Xu and Tarek Saadawi. Revealing the problems with 802.11 medium access control protocol in multi-hop wireless ad hoc networks. *Computer Networks*, 38(4):531–548, March 2002.