Quality of Service Schemes for IEEE 802.11

A Simulation Study

Anders Lindgren, Andreas Almquist, Olov Schelén
 Division of Computer Science and Communications
 Department of Computer Science and Electrical Engineering
 Luleå University of Technology, SE - 971 97 Luleå, Sweden

{Anders.Lindgren, Andreas.Almquist, Olov.Schelen}@sm.luth.se

Abstract—This paper analyzes and compares four different mechanisms for providing QoS in IEEE 802.11 wireless LANs. We have evaluated the IEEE 802.11 mode for service differentiation (PCF), Distributed Fair Scheduling, Blackburst, and a scheme proposed by Deng *et al.* using the *ns-2* simulator. The evaluation covers medium utilization, access delay, and the ability to support a large number of high priority mobile stations. Our simulations show that PCF performs badly, and that Blackburst has the best performance with regard to the above metrics. An advantage with the Deng scheme and Distributed Fair Scheduling is that they are less constrained, with regard to the characteristics of high priority traffic, than Blackburst is.

I. INTRODUCTION

As usage and deployment of wireless Local Area Networks (WLANs) increases, it is reasonable to expect that the demands to be able to run real-time applications will be the same as on wired networks. Given the relatively low bandwidth in these networks, the introduction of Quality of Service is indispensable.

The IEEE 802.11 standard [6] for WLANs is the most widely used WLAN standard today. It contains a mode for service differentiation, but that has been shown to perform badly and give poor link utilization [8]. We study and evaluate four schemes for providing QoS over IEEE 802.11 wireless LANs, the PCF mode of the IEEE 802.11 standard [6], Distributed Fair Scheduling [7], Blackburst [4], and a scheme proposed by Deng *et al.* [1].

A. IEEE 802.11

IEEE 802.11 has two different access methods, the mandatory Distributed Coordinator Function (DCF) and the optional Point Coordinator Function (PCF). The latter aims at supporting real-time traffic.

A.1 Distributed Coordinator Function

The Distributed Coordinator Function is the basic access mechanism of IEEE 802.11. It uses a Carrier Sense Multiple Access with Collision Avoidance (CSMA/CA) algorithm to mediate access to the shared medium [6].

Before sending a frame, the medium is sensed, and if it is idle for at least a DCF interframe space (DIFS), the frame is transmitted. Otherwise, a backoff time B (measured in time slots) is chosen randomly in the interval $[0, CW]$, where CW is the Contention Window. Whenever the medium has been idle for at least a DIFS, the backoff timer is decremented with one each time slot the medium remains idle. When the backoff timer reaches

zero, the frame is transmitted. If a collision is detected (which is done by the use of a positive acknowledgment scheme), the contention window is doubled and a new backoff time is chosen. The backoff mechanism is also used after a successful transmission before sending the next frame. After a successful transmission, the contention window is reset to its start value, CW_{min} .

A.2 Point Coordinator Function

PCF is a centralized, polling-based access mechanism which requires the presence of a base station that acts as Point Coordinator (PC). If PCF is to be used, time is divided into superframes where each superframe consists of a contention period where DCF is used, and a contention-free period (CFP) where PCF is used as shown in Fig. 1. The CFP is started by a beacon frame sent by the base station, using the ordinary DCF access method. Therefore, the CFP may be shortened since the base station has to contend for the medium.

Fig. 1. The IEEE 802.11 superframe

During the CFP, the PC polls each station in its polling list (the high priority stations), when they are clear to access the medium. To ensure that no DCF stations are able to interrupt this mode of operation, the interframe space (IFS) between PCF data frames is shorter than the usual IFS (DIFS). This time is called a PCF interframe space (PIFS). To prevent starvation of stations that are not allowed to send during the CFP, there must always be room for at least one maximum length frame to be sent during the contention period.

B. DENG

Deng and Chang proposes a method (which we call the DENG scheme) for service differentiation with minimal modifications of the IEEE 802.11 standard [1]. It uses two properties of IEEE 802.11 to provide differentiation: the interframe space (IFS) used between data frames, and the backoff mechanism. If two stations use different IFS, a station with shorter IFS will get higher priority than a station with a longer IFS. To further

TABLE I

DENG PRIORITY CLASSES. COMBINING BACKOFF ALGORITHMS AND IFS GIVES PRIORITIES 0-3. ρ IS A RANDOM VARIABLE IN THE INTERVAL (0, 1), AND i MEANS THE i TH BACKOFF PROCEDURE FOR THIS FRAME.

Priority	IFS	Backoff algorithm
0	DIFS	$B = \frac{2^{2+i}}{2} + \rho \times \frac{2^{2+i}}{2}$
1	DIFS	$B = \rho \times \frac{2^{2+i}}{2}$
2	PIFS	$B = \frac{2^{2+i}}{2} + \rho \times \frac{2^{2+i}}{2}$
3	PIFS	$B = \rho \times \frac{2^{2+i}}{2}$

extend the number of available classes, different backoff algorithms are used depending on the priority class. Table I shows the four defined priority classes [1].

C. Distributed Fair Scheduling

In [7] an access scheme called Distributed Fair Scheduling (DFS) which utilizes the ideas behind fair¹ queuing [2] in the wireless domain is presented. It uses the backoff mechanism of IEEE 802.11 to determine which station should send first. Before transmitting a frame, the backoff process is always initiated. The backoff interval calculated is proportional to the size of the packet to send and inversely proportional to the weight of the flow. This causes stations with low weights to generate longer backoff intervals than those with high weights, thus getting lower priority. Fairness is achieved by including the packet size in the calculation of the backoff interval, causing flows with smaller packets to get to send more often. This gives flows with equal weights the same bandwidth regardless of the packet sizes used. If a collision occurs, a new backoff interval is calculated using the backoff algorithm of the IEEE 802.11 standard.

D. Blackburst

The main goal of Blackburst [4, 5] is to minimize the delay for real-time traffic. Unlike the other schemes it imposes certain requirements on the high priority stations. Blackburst requires: 1) all high priority stations try to access the medium with equal, constant intervals, t_{sch} ; and 2) the ability to jam the medium for a period of time.

When a high priority station wants to send a frame, it senses the medium to see if it has been idle for a PIFS and then sends its frame. If the medium is busy, the station waits for the medium to be idle for a PIFS and then enters a black burst contention period. The station now sends a so called black burst to jam the channel. The length of the black burst is determined by the time the station has waited to access the medium, and is calculated as a number of *black slots*. After transmitting the black burst, the station listens to the medium for a short period of time (less than a black slot) to see if some other station is sending a longer black burst which would imply that the other station has waited longer and thus should access the medium first. If the medium is idle, the station will send its frame, otherwise it will wait until the medium becomes idle again and enter another black burst

¹Fair in the sense that each flow is allocated bandwidth proportional to some weight.

contention period. By using slotted time, and imposing a minimum frame size on real time frames, it can be guaranteed that each black burst contention period will yield a unique winner [4].

After the successful transmission of a frame, the station schedules the next transmission attempt t_{sch} seconds in the future. This has the nice effect that real-time flows will synchronize, and share the medium in a TDM fashion [4]. This means that unless some low priority traffic comes and disturbs the order, very little blackbursting will have to be done once the stations have synchronized.

Low priority stations use the ordinary CSMA/CA access method of IEEE 802.11.

II. SIMULATION SETUP

To evaluate the above described methods described, we used the network simulator *ns-2* [3] which already has IEEE 802.11 DCF functionality. We extended the simulator with implementations of IEEE 802.11 PCF and the other schemes, and ran the simulation scenarios described below to measure three different metrics: throughput, access delay and maximum number of high priority stations.

A. Scenarios

Our simulation topology consisted of several wireless stations and one base station (connected to a wired node which serves as a sink for the flows from the wireless domain) in the wireless LAN as shown in Fig. 2.

Fig. 2. An example of our simulation scenario

The traffic in our simulations was generated by each station generating constant bit rate flows to the sink. We always used 230 byte frames (including IP and UDP headers), but varied the inter-frame interval between the simulations to vary the offered load, calculated as shown in (1).

$$load = \frac{n_{stations} \cdot \frac{size_{pkt}}{interval_{pkt}}}{C_{bitrate}} \quad (1)$$

Each point in our plots is an average over ten simulation runs, and the error bars indicate the 95% confidence interval. In the

delay and throughput comparison simulations, we had 20 wireless stations and varied the fraction of high priority stations. When we investigated the maximum number of high priority stations we used a variable number of stations.

B. Metrics

The metrics we have used are *throughput*, *access delay*, and *maximum number of high priority stations*. To be able to see the differentiation and medium utilization of the schemes, we have looked at both the average throughput for the stations at each priority level, and the total throughput for all stations together. To compare the graphs from different levels of load, we plot a normalized throughput on the y axis which is calculated as the fraction of the offered data actually delivered to the destination.

To determine to what extent the schemes were able to provide good service to high priority traffic, we ran simulations where the stations sent 65.7 kbit/s streams. We fixed the low priority traffic load at certain levels, and gradually increased the number of high priority stations to see how many simultaneous high priority stations that could get good service. We used two definitions of good service. The first considers throughput, and requires that 95% of the offered data is delivered, while the second requires access delay to be below 20 ms.

C. Method specific details

Table II shows the parameter values used in our simulations. For further explanation and description of the parameters, we refer to [1, 4, 6, 7].

TABLE II
PARAMETER VALUES USED IN OUR SIMULATIONS

Parameter	Value	Parameter	Value
DIFS	50 μ s	Time slot	20 μ s
PIFS	30 μ s	$C_{bitrate}$	2 Mbit/s
Superframe	110 TU^2	$size_{pkt}$	230 bytes
Max CFP	108.85 TU	CW_{min}	31
Parameter	Value	Parameter	Value
Deng high prio	3	DFS high weight	0.075
Deng low prio	1	DFS low weight	0.025
Deng DIFS	100 μ s	DFS Scaling_Factor	0.02
Black slot	20 μ s		

When using PCF, during a CFP the Point Coordinator (the base station) polls the stations in its polling list in a round robin fashion. If all stations have been polled once, the CFP will be ended prematurely. If there is not enough time to poll all stations the next station in the list will be polled first in the next CFP. To enhance the performance of DFS when there is much low priority traffic, we decided to use exponential mapping [7] of the backoff intervals.

III. RESULTS

Our initial simulations compared the performance of the different schemes with regard to *throughput*. The simulations show that even at low loads, PCF gives low priority flows significantly

lower throughput than the other schemes do. The PCF high priority stations perform acceptable at this low load, but the performance for these starts to deteriorate when the amount of high priority traffic increases. Fig. 3 shows how well the different schemes provides service differentiation with regard to throughput, and Fig. 4 shows the total throughput, which indicates how well the different schemes utilizes the medium. We have run simulations with several levels of load but because of space limitations we only present the most interesting graphs here.

As we increased the load in our simulations (see Fig. 3, the right graph), none of the schemes were capable of delivering all data of the high priority stations when there are only high priority stations in the system. Blackburst gives the best performance both for high and low priority traffic. This also implies that Blackburst has the best medium utilization, verified in Fig. 4.

An interesting observation is that the throughput for low priority traffic cases increases slightly for PCF and DENG when there is only one low priority station. Our hypothesis about this is that all high priority stations will send their frames in what appears to the low priority stations as a big “chunk” (not letting any low priority traffic get in between their frames). After that, all high priority stations will start decrementing their backoff timers, not contending for the medium. During this time, low priority stations can access the medium. When there is only one low priority station, it will get to send, without contending with some other low priority station. A similar phenomenon occurs for Blackburst.

When investigating the second metric, *access delay*, we found that Blackburst and Deng performs well for high priority traffic. Blackburst also gives low access delay to low priority traffic as long as the load is relatively low, but it should be noted that when the network becomes heavily loaded Blackburst totally starves low priority traffic. As shown in Fig. 5 one can see that the access delay increases as the fraction of high priority traffic increases for all schemes.

The investigation of the third metric, *maximum number of high priority stations*, indicates that Blackburst is the scheme capable of supporting the largest number of prioritized stations both with regard to throughput and access delay. As Fig. 6 shows, both Blackburst and DENG are able to give the high priority stations good service, regardless of the amount of low priority traffic. The reason why DFS doesn’t perform that well is due to the fact that DFS tries to distribute the bandwidth fairly among the stations according to their weights instead of trying to give perfect service to high priority traffic. For PCF this is both because it has poor medium utilization, and because there must always be room for a low priority frame during the contention period.

In a real life scenario, it is not likely that all traffic is CBR and therefore we also ran simulations with burstier traffic³ which did not affect the high priority traffic in any significant way. Thus we feel that the results presented here would be valid even with other characteristics of low priority traffic.

³ON-OFF sources with ON and OFF periods from a Pareto distribution

²1 $TU = 1024 \mu$ s

Fig. 3. Average throughput for a station at the given priority level.

Fig. 4. Total throughput for the QoS schemes.

Fig. 5. Average access delay for a station at each priority level.

IV. CONCLUSIONS

From our simulations we can conclude that the PCF mode of the IEEE 802.11 standard performs poorly in the metrics stud-

ied compared to the other schemes evaluated. Blackburst gives the best performance to high priority traffic both with regard to throughput and access delay. A drawback with Blackburst is the requirements it imposes on the high priority traffic. If these

Fig. 6. Maximum number of high priority stations with good performance.

can not be met, DENG might be a suitable alternative since it can serve quite many high priority stations, while giving them very low access delay. An major advantage of DFS is that it will try to achieve fairness, and will not starve low priority traffic, which in many cases is a desirable property of a scheme. Further, our simulations show that Blackburst is the scheme among those studied here that gives the best medium utilization, which is important, given the scarcity of bandwidth in wireless networks.

Finally, we conclude with the observation that there might not be one scheme that is the best to choose in all situations, but the choice of QoS scheme should instead depend on the expectations of the traffic, and other circumstances. Before deciding on what QoS scheme to use in a network, an analysis of what the network should be used for, and what kind of services that is needed should be done.

ACKNOWLEDGMENTS

First of all we would like to thank João Sobrinho for helping us to gain a better understanding of Blackburst. We would also like to express our deepest gratitude to Nitin H. Vaidya, Paramvir Bahl, and Seema Gupta for their benevolence in providing us with their *ns* implementation of DFS, saving us many long hours of implementation work.

REFERENCES

- [1] D-J. Deng and R-S. Chang. A priority scheme for IEEE 802.11 DCF access method. *IEICE Transactions on Communications*, E82-B(1), January 1999.
- [2] S. J. Golestani. A self-clocked fair queueing scheme for broadband applications. In *Proceedings of IEEE INFOCOM*, 1994.
- [3] S. McCanne and S. Floyd. *ns network simulator version 2.1b6*, August 2000.
- [4] J. L. Sobrinho and A. S. Krishnakumar. Real-time traffic over the IEEE 802.11 medium access control layer. *Bell Labs Technical Journal*, pages 172–187, Autumn 1996.
- [5] J. L. Sobrinho and A. S. Krishnakumar. Quality-of-Service in ad hoc carrier sense multiple access networks. *IEEE Journal on Selected Areas in Communications*, 17(8):1353–1368, August 1999.
- [6] The Institute of Electrical and Electronics Engineers, Inc. *IEEE Std 802.11 - Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) specifications*, 1999 edition.
- [7] N. H. Vaidya, P. Bahl, and S. Gupta. Distributed fair scheduling in a wireless LAN. In *Sixth Annual International Conference on Mobile Computing and Networking*, Boston, August 2000.

- [8] M. A. Visser and M. El Zarki. Voice and data transmission over an 802.11 wireless network. In *Proceedings of PIMRC'95, Toronto, Canada*, pages 648–652, September 1995.