Evaluation of Quality of Service Schemes for IEEE 802.11 Wireless LANs

Anders Lindgren, Andreas Almquist, Olov Schelén
Division of Computer Science and Networking
Department of Computer Science and Electrical Engineering
Luleå University of Technology, SE - 971 87 Luleå, Sweden
{Anders.Lindgren, Andreas.Almquist, Olov.Schelen}@sm.luth.se

Abstract

This paper evaluates four mechanisms for providing service differentiation in IEEE 802.11 wireless LANs, the Point Coordinator Function (PCF) of IEEE 802.11, the Enhanced Distributed Coordinator Function (EDCF) of the proposed IEEE 802.11e extension to IEEE 802.11, Distributed Fair Scheduling (DFS), and Blackburst using the ns-2 simulator. The metrics used in the evaluation are throughput, medium utilization, collision rate, average access delay, and delay distribution for a variable load of real time and background traffic. PCF performance is comparably low, while EDCF performs much better. The best performance is achieved by Blackburst. DFS provides relative differentiation and consequently avoids starvation of low priority traffic.

1. Introduction

The IEEE 802.11 standard [5] for WLANs is the most widely used WLAN standard today. It has a mode of operation that can be used to provide service differentiation, but it has been shown to perform badly [7]. We study and evaluate four schemes for providing QoS over IEEE 802.11 wireless LANs; the PCF mode of the IEEE 802.11 standard [5], Distributed Fair Scheduling [6], Blackburst [4], and Enhanced DCF [1]. This paper is a continuation of previous work where some initial comparisons between QoS schemes were done [3]. This paper contains more realistic traffic scenarios, and some new metrics. Furthermore, this paper also evaluates the EDCF access mechanism of the upcoming IEEE 802.11e standard.

2. Overview of evaluated schemes

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.

DCF is the basic access mechanism of IEEE 802.11, and uses a Carrier Sense Multiple Access with Collision Avoidance (CSMA/CA) algorithm to mediate the access to the shared medium. Before a data frame is sent, the station senses the medium. If it is idle for at least a DCF interframe space¹ (DIFS) period of time, 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 so called Contention Window, calculated as $CW_i = 2^{k+i-1} - 1$, where i is the number of attempts (including the current one) to transmit the frame that has been done, and k is a constant defining the minimum contention window, CW_{min} . After the medium has been detected idle for at least a DIFS, the backoff timer is decremented by one for each time slot the medium remains idle. When the backoff timer reaches zero, the frame is transmitted. Upon detection of a collision, a new backoff time is chosen and the backoff procedure starts over. Because the contention window is exponentially increased, the risk of further collisions is reduced. 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 CW_{min} .

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 supported, both PCF and DCF coexist and in this case, time is divided into superframes. Each superframe consists of a contention period where DCF is used, and a contention free period (CFP) where PCF is used. During the CFP, it sends poll frames to 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 IFS between PCF data frames is shorter than the usual DIFS. This space is called a PCF interframe space (PIFS). To prevent starvation of low

¹ An interframe space, IFS, is the time a station waits when the medium is idle before attempting to access it. IEEE 802.11 defines several IFSs, and by using shorter IFS, the medium is accessed prior to stations using a longer IFS. This is e.g. used to ensure that an acknowledgment frame is sent before any other station can send data.

priority flows, the contention period must always be long enough for one maximum length frame.

IEEE 802.11e – Enhanced DCF Task group E of the IEEE 802.11 working group are currently working on an extension to the IEEE 802.11 standard called IEEE 802.11e. The goal of this extension is to enhance the access mechanisms of IEEE 802.11 and provide a distributed access mechanism that can provide service differentiation. All the details have not yet been finalized, but a new access mechanism called Enhanced DCF (EDCF), which is an extension of the basic DCF mechanism, has been selected [1].

EDCF combines two measures to provide differentiation. The minimum contention window (CW_{min}) can be set differently for different priority classes, yielding higher priority to classes with smaller CW_{min} . For further differentiation, different interframe spaces can be used by different traffic classes. Instead of DIFS, an interframe space called Arbitration Interframe Space (AIFS) is used. The AIFS for a given class should be a DIFS plus some (possibly zero) time slots. Classes with smaller AIFS will have higher priority.

To enhance the performance, and achieve better medium utilization, *packet bursting* can be used [2], meaning that once a station has gained access to the medium, it can be allowed to send more than one frame without contending for the medium again. After getting access to the medium the station is allowed to send as many frames it wishes as long as the total access time does not exceed a certain limit (TxOpLimit). To ensure that no other station interrupts the packet burst, a shorter IFS than usual is used between packets. If a collision occurs, the packet burst is terminated. Since packet bursting might increase the jitter, TxOpLimit should not be longer than the time required for the transmission of a data frame of maximum size.

Distributed Fair Scheduling Vaidya et al. proposes an access scheme which utilizes the ideas behind fair queuing in the wireless domain, called Distributed Fair Scheduling (DFS) [6].

DFS uses the backoff mechanism of IEEE 802.11 to determine which station should send first. The backoff interval will be longer the lower the weight of the sending station is, so differentiation will be achieved, while fairness is achieved by making the interval proportional to the packet size.

Blackburst Sobrinho and Krishnakumar proposes a scheme called Blackburst, with the main goal of minimizing delay for real time traffic [4]. Blackburst requires that all high priority stations try to access the medium with constant intervals, t_{sch} . Further, Blackburst also requires the ability to jam the wireless medium for a period of time. Low priority stations use the ordinary DCF access mechanism of IEEE 802.11.

If the medium is found busy when a station wants to transmit real-time data, the station waits until it becomes idle and then enters a black burst contention period by jamming the channel for a period of time. The length of the black burst is determined by the time the station has been waiting to access the medium. After transmitting the black burst, the station listens to the medium to see if some other station is sending a longer black burst, implying 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 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 access instant (when the station will try to transmit the next frame) t_{sch} seconds in the future. By doing this, real-time flows will synchronize, and share the medium without collisions, meaning that very little black-bursting will have to be done once the stations have synchronized [4].

3. Simulations

To evaluate the methods described above, we use the simulator *ns-2*. Our simulations consist of traffic that has been chosen to be similar to data generated by for example a variable bit rate audio or video encoder, and some low priority background traffic. Each wireless station initiates a flow to a sink located beyond the base station of the wireless LAN. The high priority stations generate packets with packet sizes taken from a normal distribution with mean 300 bytes, and standard deviation 40 bytes. We have used inter-packet intervals of 25 and 40 ms, which gives us data flows with an average bit rate of 96 and 60 kbit/s. The low priority stations generate packets every 50 ms, with a packet size taken from a normal distribution with mean 800 bytes, and standard deviation 150 bytes (corresponding to a bit rate of 128 kbit/s). We have had some fixed numbers of low priority stations (3 and 12 stations), and gradually increased the number of high priority stations to increase the load of the system. All wireless stations are located such that every station is able to detect a transmission from any other station, and there is no mobility in the system².

Metrics The *average throughput* for the stations at each priority level, shows how well the QoS schemes can provide service differentiation between the various priority levels. To be able to compare the graphs from different levels

²Further simulation details, including parameter settings, and graphs omitted because of space limitations, can be found at <http://www.sm.luth.se/~dugdale/publications>

Figure 1. Comparison of schemes with regard to throughput, medium utilization, and access delay.

of load, we plot a normalized throughput, calculated as the percentage of the offered data that is actually delivered to the destination. Because of the scarcity of wireless bandwidth, we also study the *medium utilization* of the different schemes, by measuring how large percentage of time that is used for transmission of data frames. The *collision rate* is the average number of collisions that occur per second. *Access delay* is the time from when a packet reaches the MAC layer until it is successfully transmitted. We measure the average access delay to see how well the schemes can accommodate real-time flows. However, for real-time flows it is often not enough with a low average access delay, but there can be delay bounds after which the data is useless. We present the *cumulative distribution of access delays* for high priority traffic to find out the percentage of packets that are below certain delay bounds.

4. Results

In Fig. 1(a), we can see that the Blackburst scheme gives the best performance to high priority traffic with regard to throughput, especially at lower loads. However, at higher loads we see that EDCF also has very good performance for high priority traffic, while it starts to deteriorate rather early for PCF, and somewhat after that for DFS as well. On the other hand, these schemes give better performance to low priority traffic while Blackburst and EDCF completely starves it at high loads. Fig. 1(b) shows the intuitive result

that Blackburst has the best medium utilization as well. The use of packet bursting for EDCF makes it reasonable to believe that utilization will be high (since less contention have to be done), so the rather low utilization of EDCF is surprising. One explanation to this can be found in Fig. 3, where collision rates are shown. EDCF has higher collision rates than the other schemes, which impacts the performance. Interesting to see here is that the collision rate for Blackburst decreases as the number of high priority stations increases. This verifies that no collisions occur between Blackburst nodes (all the collisions seen here are between low priority stations – something that can be seen by the fact that at the same point as low priority traffic is starved in Fig. 1(a), the collision rate reaches zero).

Looking at the average delay in Fig. 1(c), we see that all the schemes give rather acceptable average delays to the high priority traffic, but Blackburst (and to some extent EDCF) has very low delays. It is however more interesting to study the distribution of the delay in Fig. 2. Here, the impact of the

Figure 3. Collision rates.

Figure 2. Cumulative delay distribution.

packet bursting of EDCF can clearly be seen, as a large part of the packets have very low delays (the packets within a packet burst), and the packets that have to contend for the medium have longer delays. Blackburst and DFS have rather steep curves (which indicates a low variance, and little jitter) which reaches the vicinity of 100% quickly, meaning that the delay has a rather low upper bound. On the other hand, traffic using PCF and EDCF have flatter curves, and especially at high loads, parts of the packets have really high delays (over 100 ms). Since real-time applications often have a bounded tolerable delay, it can be more important that a large part of the packets have delays below that bound than to have a really low average delay. For example, assume that the maximum tolerable delay is 100 ms. At the highest load we can now see that while DFS and Blackburst manages to give a delay below that to virtually all packets, EDCF and PCF only does that to about 85% of the packets. This means that even though DFS has lower throughput for high priority traffic than EDCF, they actually deliver approximately the same amount of useful real-time data (which speaks in favor of DFS since it has better performance for low priority traffic).

5. Conclusions

Our simulations show that the new EDCF mechanism developed by the IEEE 802.11e task group is an improvement over PCF that is shown to have rather poor performance. EDCF is completely distributed, has better performance than PCF, and is less complex.

Blackburst gives the best performance to high priority traffic both with regard to throughput and access delay. At low loads, it also gives rather good performance to low priority traffic, but at high loads, low priority traffic is starved. Further, our simulations show that the Blackburst scheme gives the best medium utilization. This is important, given the scarcity of bandwidth in wireless networks. We have also verified that Blackburst avoids collisions between high priority stations. A drawback with Blackburst is the requirements of constant access intervals it imposes on high priority traffic. If these requirements can not be met, EDCF

might be a suitable alternative. Although not being able to provide as good service as Blackburst, and suffering from a high rate of collisions, it still provides good service differentiation, and give low average delay to high priority traffic (unfortunately, the distribution of delays is however such that at high loads, a rather large fraction of the packets have very long delays, which might render them useless to real-time applications). At higher loads, low priority traffic suffers from starvation just like when using Blackburst. In many cases it is not desirable to starve low priority traffic, but rather to give a relative differentiation. DFS ensures better service to high priority traffic, and still does not starve low priority traffic, but ensures that it gets its fair share of the bandwidth.

References

- [1] M. Benveniste, G. Chesson, M. Hoeben, A. Singla, H. Teunissen, and M. Wentink. EDCF proposed draft text. IEEE working document 802.11-01/131r1, March 2001.
- [2] G. Chesson, W. Diepstraten, D. Kitchin, H. Teunissen, and M. Wentink. Baseline D-QoS proposal. IEEE 802.11 working document 802.11-00/399, November 2000.
- [3] A. Lindgren, A. Almquist, and O. Schelén. Quality of Service schemes for IEEE 802.11 - a simulation study. In *Proceedings of the Ninth International Workshop on Quality of Service (IWQoS'2001)*, June 2001.
- [4] 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.
- [5] The Institute of Electrical and Electronics Engineers, Inc. *IEEE Std 802.11 - Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) specifications*. The Institute of Electrical and Electronics Engineers, Inc., 1999 edition.
- [6] 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.
- [7] 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.