

A Short Look on Power Saving Mechanisms in the Wireless LAN Standard Draft IEEE 802.11

*Christian Röhl, Hagen Woesner, Adam Wolisz**
Technical University Berlin
Telecommunication Networks Group
{roehl, woesner, wolisz}@ee.tu-berlin.de

** also with GMD FOKUS*

Abstract

This paper describes simulations of the power saving mechanism of the upcoming standard for wireless Local Area Networks IEEE 802.11[1]. They were performed in order to see how typical parameters influence the performance. Simulations were made for a ad-hoc-network with 8 stations. Figures for optimum Beacon intervals and ATIM window sizes were obtained.

1. Introduction

Wireless Local Area Networks (WLANs) are a rapidly growing area in networking. This is basically due to the upcoming of portable devices like notebooks and mobile phones. A key feature of these devices is that the limited battery capacity, which limits their time in action. This results in a need of power saving mechanisms, which prolong the life time of the batteries.

The next chapter describes in short different ways to address the power saving problem. Chapter 2 shows the way power saving is implemented in the IEEE standard 802.11. After that we describe the simulated environment, the source model and the parameter set used for the simulations. In chapter 4 we come to the simulations and their results. A discussion of the results and conclusions is shown in chapter 5.

2. Power Saving in the IEEE 802.11 draft standard

In general, the best way to save power for wireless communication devices would be to switch them off. Unfortunately, one can not do this without losing the capability to communicate in both directions, i.e. a station in this kind of a power saving mode would not know of any packets arriving for it at this time. Therefore there are two problems to be addressed in power saving:

- How does a station in power save mode receive packets from other stations?
- How does a station send to another station in power save mode?

Within the standard, the general idea is for all stations in PS mode to be synchronized to wake up at the same time. At this time there starts a window in which the sender announces buffered frames for the receiver. A station that received such an announcement frame stays awake until the frame was delivered. This is easy to be done in infrastructure networks, where there is a central access point, which is able to store the packets for stations in doze state and to synchronize all mobile stations. It is more difficult for ad-hoc net-

works, where the packet store and forward and the timing synchronization has to be done in a distributed manner.

Power Saving in IEEE 802.11 therefore consists of a Timing Synchronization Function (TSF) and the actual power saving mechanism. The TSF for an infrastructure network (the Point Coordination Function - PCF) can be seen in Figure 1. The access point (AP) is responsible for generating beacons, which contain a valid time stamp beside other information. Stations within the BSS (Basic Service Set - a wireless cell) adjust their local timers to that time stamp. If the channel is in use after the beacon interval the AP has to defer its transmission until the channel is free again.

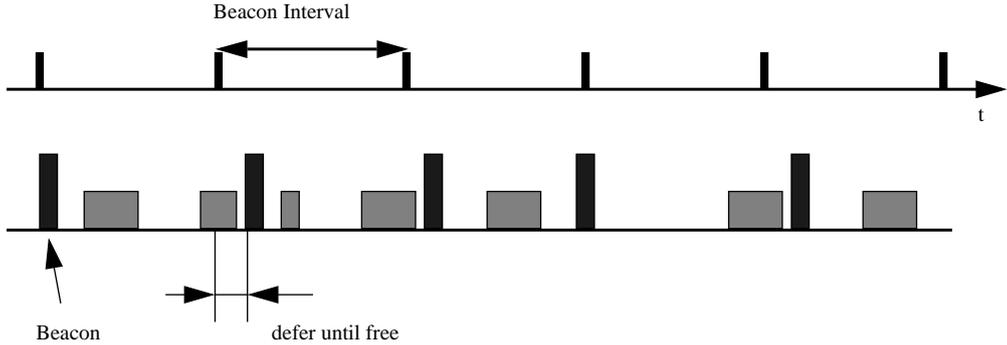


Figure 1: TSF for infrastructure networks in 802.11

The situation is more complicated for an ad-hoc network (the Distributed Coordination Function - DCF, see Figure 2). Due to the absence of a trusted authority the timers adjust in a distributed way: Every station is responsible for generating a beacon. After the beacon interval all stations compete for transmission of the beacon using the standard back-off algorithm. The first station “wins” the competition and all others have to cancel their beacon transmission and to adjust their local timers to the time stamp of the winning beacon.

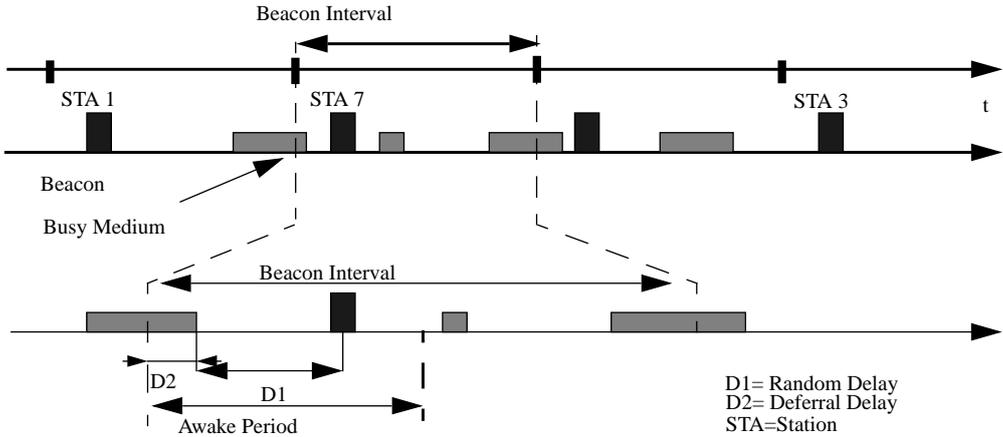


Figure 2: TSF for ad-hoc networks in 802.11

The power management in the PCF is simple due to the existence of the AP as central buffer for all packets to stations in doze mode. The AP transmits together with the beacon

a so-called Traffic Indication Map (TIM). All unicast packets for stations in doze mode are announced in the TIM. The mobiles afterwards poll the AP for the packets. If broadcast/multicast frames are to be transmitted, they are announced by a Delivery TIM (DTIM) and sent immediately afterwards. Of course the stations in power save mode have to wake up short before the end of the beacon interval and to stay awake until the beacon transmission is over.

The power management for the DCF is based on the same distributed fashion as it is used for the TSF. Packets for a station in doze state have to be buffered by the sender until the end of the beacon interval. They have to be announced using Ad-hoc TIMs (ATIMs), which are transmitted in a special interval (the ATIM window) directly after the beacon. ATIMs are unicast frames which have to be acknowledged by the receiver. After sending the acknowledgment, the receiver does not fall back into doze state but stays awake and waits for the announced packet (see Figure 3). Both ATIMs and the data packets have to be transmitted using the standard backoff algorithm.

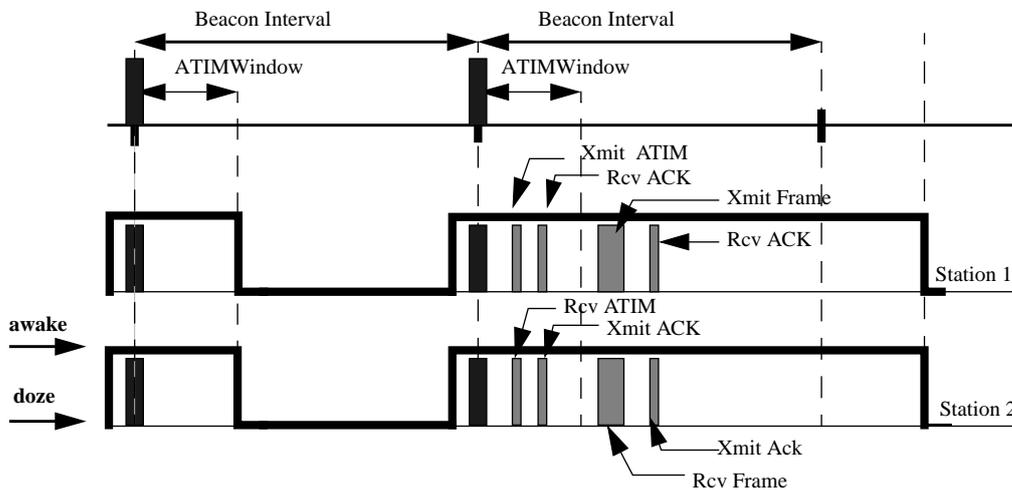


Figure 3: Power Management in the DCF of 802.11

Our aim was to tune the algorithm to get best values for the throughput of stations in power save mode and on the other hand for a maximum possible time in doze state. We chose the ratio of time in doze state vs. the time in active state as a measure for the quality of the power saving mechanism itself.

3. Simulation Approach

Our simulations were performed using a PTOLEMY model [2], which is described in greater detail in [3][3]. We used the appropriate values for the Direct Sequence Spread Spectrum (DSSS) physical layer. The simulation environment consists of 8 stations, which belonged to an independent basic service set (IBSS). We did not consider any hidden terminals. Simulations with 1, 2, 4 and all 8 stations in power save mode were performed.

To model realistic traffic on a wireless LAN, we used trace files of an Ethernet [4], which were multiplexed using different start points in the file to lead to different, but predictable load scenarios¹. We simulated overall offered loads of around 15, 30 and 60% of the

raw physical throughput.

We made the assumption that power consumption is proportional to the time in active mode. Any additional effects which are depending on the PHY layer like equalization and on-off switching costs were not taken into account.

4. Simulation Results

First we wanted to observe dependencies of the throughput compared to different window sizes for the beacon interval and ATIM window. As it may be expected, higher numbers of stations in power save mode lead to lower throughput. This is because of the overhead for each data packet, which consists of an ATIM and an ACK and two backoff sequences, regardless of the size of the packet to be transmitted.² It showed that there is a decrease in throughput for very small and very large ATIM window sizes (see Figure 4). An ATIM window which is too small results in less ATIMs and therefore in less packets, which can be announced and transmitted. On the other hand, when the ATIM window is too large, more ATIMs are sent than there is actually time for the packets.

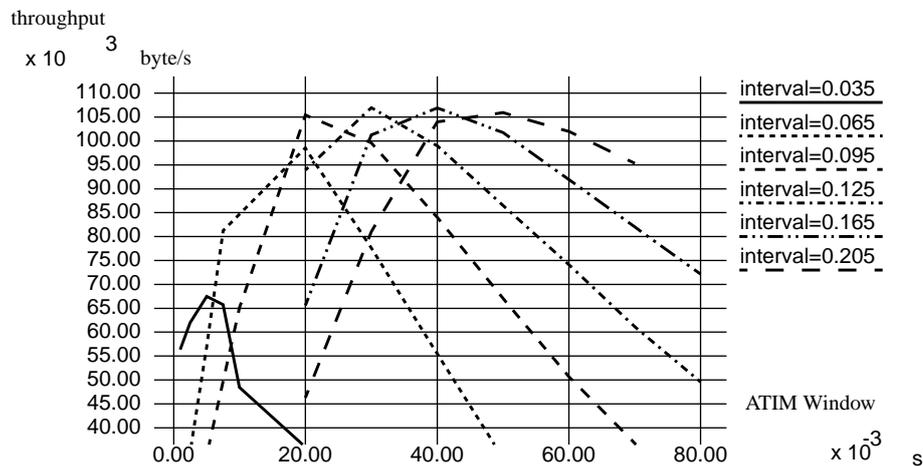


Figure 4: Throughput vs. ATIM window size for different beacon intervals, load=60.76%, 8 stations in power save mode

When we used a lower offered load for the simulations the results were basically the same, though throughput was constant for a broader range of ATIM window sizes. This was due to the fact that the channel could not be saturated any more.

In result it was obvious that the ATIM window size should be proportional to the beacon interval and that it should take 1/4 to 1/3 of the beacon interval.

The next question was to determine the time in doze state in relation to the total time. In Figure 5 one can see that the time in doze state increases when using shorter beacon intervals.

1. The trace files were recorded on our institute-internal 10Base2 Ethernet
2. A comparable overhead applies for the optional RTS/CTS exchange, which in contrast depends on the packet length.

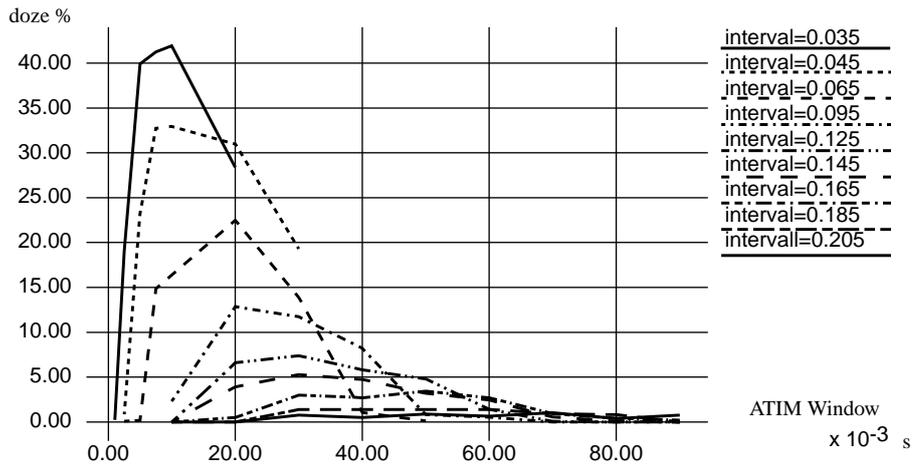


Figure 5: Percentage of time in doze state vs. ATIM window size for different beacon intervals, load=30.72%, 8 stations in PS mode

The simulation shown here was performed at an offered load of about 30%, because at a higher offered load a station would probably not fall into doze state very often.

In Figure 6 we simulated the same scenario as before, but with an offered load of about 15%. It shows that a station can stay in doze mode up to 70% of the time for beacon intervals small enough to allow for a fast transmission of the packet.

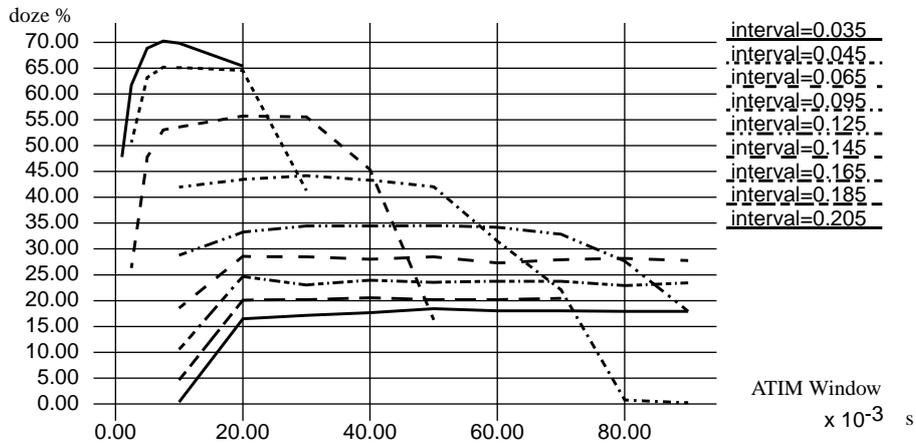


Figure 6: Percentage of time in doze state vs. ATIM window size for different beacon intervals, load =15%

The results can be explained as follows: The bigger the beacon interval the bigger the possibility that a station wishes to send during that time. This means that it has to transmit ATIMs in almost every beacon interval and to stay awake until the transmission is completed. The same applies for a receiving station. In addition to that, more ATIMs per beacon interval have to be transmitted in bigger beacon intervals, which leads to a higher collision rates and longer medium access times.

5. Conclusions

In this paper we presented simulations of the power saving mechanism in ad-hoc networks using the IEEE 802.11 standard. Work on this simulations started at a time when there was no recommendation for certain values of the parameters in the current version of the draft standard. In the meantime the values of interest in this scope are set to be 100 ms for the beacon interval and only 4 ms for the ATIM window.

Based on this work we can recommend figures for the ATIM window and beacon interval. Generally the mechanism gets less sensitive against the ATIM window size with higher values for the beacon interval. The simulations showed an optimum for the throughput at about 95 ms beacon interval. The ratio between ATIM window and beacon interval should be 1/4 to 1/3. While the first result corresponds to the value in the draft quite well, there is a „slight“ difference in our recommended size for the ATIM window. This should be explained as follows: The recommended ATIM window size of only 4 ms (or $K\mu s$, according to the standard) will be too small if there are many stations in power save mode or if the overall load is above 10%. We would definitely recommend for a higher value of the ATIM window parameter. On the other hand, there should be a means to adapt the value of this parameter to the offered load or, to be more exact, to the sum of the offered loads of the stations in PS mode.

The beacon interval should be smaller to lead to longer times in doze state. There has to be a trade-off between power saving and the overhead needed for it. If we would sacrifice about 10% in throughput we can save up to 30% more energy.

References

- [1] The editors of IEEE: Wireless LAN Media Access Control (MAC) and Physical Layer (PHY) Specification, IEEE 802.11 Draft Version 4.0; May 1996
- [2] PTOLEMY; www-site: <http://www.ptolemy.eecs.berkeley.edu>, Copyright 1990-96 The Regents of the University of California
- [3] J. Weinmiller, H. Woesner, A. Wolisz: „Analyzing and Tuning the Distributed Coordination Function in the IEEE 802.11 Draft Standard“; June 1995 MASCOT '96; Feb. 1996, San Jose, California
- [4] W. Leland et al. :“On the self-similar nature of Ethernet traffic“; IEEE Transactions on Networking, Vol. 2, No. 1, Feb.'94
- [5] R. Caceres, V. Padmababhan: „Fast and Scalable Handoffs for Wireless Internetworks“; Proc. of MobiCom, November 1996