Reducing Power Consumption in IEEE 802.11 Networks

Jean Lorchat
LSIIT - CNRS
University Louis Pasteur
Pole API - Bd Sebastien Brant
67400 Illkirch - France
Email: lorchat@dpt-info.u-strasbg.fr
Telephone: (+33) 3 902 445 87
Fax: (+33) 3 902 444 55

Thomas Noel
LSIIT - CNRS
University Louis Pasteur
Pole API - Bd Sebastien Brant
67400 Illkirch - France
Email: noel@dpt-info.u-strasbg.fr
Telephone: (+33) 3 902 445 92
Fax: (+33) 3 902 444 55

Abstract **In this paper, we show a performance evaluation of the frame aggregation mechanism. This mechanism is introduced as a replacement for the Power Save Polling (PSP) mode integrated in the IEEE 802.11 standard [1]. Since it uses bigger frame to achieve lower traffic and overhead, we investigate the consequences of packet loss on power consumption. Using this study, we are able to propose an adaptive mechanism to trigger frame aggregation based on the medium conditions as they can be observed from the station (STA) standpoint in an infrastructure environment.**

I. INTRODUCTION

Wireless communications have experienced a major growth worldwide in the past decades pushed by technologies that have been massively endorsed by the public among which the most well known are GSM for cellular telephony networks, IEEE 802.11 for Wireless LANs and Bluetooth for Wireless PANs.

Following these trends, new innovative devices have shown up to take full advantage of wireless communications. These devices provide the functionality of a computer with highly reduced encumbrance. The design of such devices was made possible thanks to the miniaturization of electronic circuits and these are becoming invaluable to more and more people in the shape of Personal Digital Assistants, Smartphones, and so on.

However these new wireless devices designed to use the wireless access technologies are battery powered and require power aware protocols to perform at their best, without draining tremendous amounts of power for communications. The power consumption problem has been studied prior to the standardization of the IEEE 802.11 protocol and it was thus made power efficient by adding an optionnal power saving mode called Power Save Polling (known as PSP). However this mode heavily impairs realtime communications because it imposes a large additional latency on each packet and leads to degraded performance both in terms of latency and throughput [2].

There is a crucial need for energy efficient protocols in Wireless LANs because hotspots are getting widespread and battery powered devices need to communicate without jeopardizing the battery lifetime. From then on there are two ways to proceed to save battery power on handhelds devices that communicate using a WLAN network interface : replace the existing PHY/MAC layers pair, or modify them to be more power efficient. The first proposition is out of the scope of this article because such surveys already exist that compare new MAC protocols for wireless LANs [3]. We will especially focus on the second part of the alternative in this paper.

In the next section, we present a brief survey of the existing PSP mode from the IEEE 802.11 standard, and modifications that were proposed for the IEEE 802.11 MAC to improve the energy efficiency. We then present the simulation environment and methodology we set up in order to evaluate the power performance of a particular mechanism. Afterwards we present and discuss the results obtained and use them to define a model for adaptive behavior of the Frame Aggregation mechanism.

II. POWER CONSUMPTION IN IEEE 802.11 WLAN CELLS

A. State of the Art

As we already stated in previous section, the IEEE 802.11 standard provides an energy conservation mechanism named PSP. This mechanism relies on periodic signals transmitted by the Access Point (AP) in infrastructure networks : the beacon. These beacons have a fixed period and stations wanting to save power with PSP can sleep between successive beacons, provided they negotiated entering PSP mode with their AP. The beacon is then tagged when traffic for sleeping stations is pending, and stations can pick their pending messages up right after the tagged beacon. This mode allows to achieve very high energy savings because the beacon interval suggested by the standard is 100 milliseconds, and stations that do not communicate spend most of the time sleeping but stay reachable within the time of one beacon interval. Though as soon as the station begins to communicate, PSP mode has to be stopped unless latency is not an issue. This alternation of PSP and Constantly Awake Mode (CAM) is called FastPSP or PSPCAM and was proposed in implementations (i.e. it is not required by the standard, but rather available as an extension to this standard).

0-7803-9182-9/05/$20.00 ©2005 IEEE.

The issue is the following : using PSP we have very good results but bad latency. Furthermore, PSP is only saving energy for incoming communications because outgoing frames trigger the wakeup of the station. If we want to have little impact on latency, we can use FastPSP but then there is no power saving mechanism used when the station is not idle. We aim to improve the power efficiency of the IEEE 802.11 MAC protocol by providing extensions that are backwards compatible. This means that stations without support for new extensions will bypass the messages without understanding them, but will not be affected in any other way.

B. Our Proposed Mechanism

The mechanism we proposed in [4] is called *Frame Aggregation*. It takes advantage of the higher Maximum Transmission Unit (MTU) in WLAN cells to collect successive small frames and put them together in a bigger frame. The backward compatibility is ensured by using new MAC level frame types that will be discarded by legacy stations. The AP and the station negotiate the ability to use Frame Aggregation by means of new optional elements in Beacons (for the AP) and Association Request (for the station) messages.

By putting several frames together, we save the overhead caused by lower layers for every frame but one. However we have to investigate the impact of this mechanism on packet latency because the mechanism waits for frames to be ready for transmission, up to the temporal aggregation threshold limit. In addition, the wireless nature of the communications and the associated collisions or errors that occur have to be taken into account.

III. PERFORMANCE EVALUATION METHODOLOGY

The most important parameter that can impede energy savings using Frame Aggregation is frame loss that can occur because of bit errors on the link as well as collisions. When a frame is lost, a retry takes place at MAC level in both the regular case and using Frame Aggregation. However, the aggregated frame is larger and will consume more power to be retransmitted than the regular frame. We have to investigate the point (if exist) where Frame Aggregation becomes less power efficient than regular mode.

A. Cell Simulation

We will conduct this analysis using simulations of an IEEE 802.11 cell (also known as a Basic Service Set : BSS). The simulator is a simple discreet event based simulator [5] implementing basic functionalities of the cell from the IEEE 802.11b standard.

The key additional features we require are :

- latency characterization which allows us to know the latency a frame has suffered from the time it entered the First In First Out (FIFO) queue in the transmitting station up to the time it arrived without errors at the receiving station.
- a power consumption model derived from preliminary experimentations. We use experimental values that we obtained from a working wireless card using power dissipation measurements (much like in [6] and [7]).
- a proper Frame Aggregation implementation that respects the behavior of the latency characterization property so as to compare efficiency of the Frame Aggregation mechanism with the legacy IEEE 802.11 operating mode.

In addition, we want to be able to adjust the packet error rate (PER) occuring within the cell in order to discover key points where Frame Aggregation would be less efficient.

B. Scenario

The scenario we will operate comprises both fixed parameters and variable parameters. The variable parts will change from one simulation run to another, but all parameters are fixed within a given simulation run.

We first introduce the variable parts. The number of station in the cell will range from 1 to 7 (not including the AP). The PER will range from 10^{-1} to 10^{-5} in one hundred times increments. The flow will first be directed from the station to the AP, and then from the AP to the station. Of course, we will run simulations for regular mode and Frame Aggregation.

The fixed part is mainly the flow that we will use. It is a constant bitrate flow consisting of 512 bytes level 3 packets. The delay between each frame is four milliseconds which amounts to about 128 kilobytes per second, big enough to carry high quality audio or low quality video.

Looking at this flow, we can anticipate that Frame Aggregation mode will build Frame Aggregation superframes containing four frames, inducing at most a sixteen milliseconds additional latency for the first frame.

The simulations will run for thirty simulated minutes, and each given parameter set will be run five times.

IV. SIMULATIONS RESULTS

The results collected from simulation have been summarized in two categories : power consumption and latency performance. Whereas the first set of results allows to estimate the power efficiency of Frame Aggregation (as compared to regular mode), second is determinant for mechanism validity : should latency results be poor, our mechanism would be useless (much like PSP).

A. Power Consumption Results

In the following figures, we use a similar way to represent results for each case of the original scenario. We plotted energy consumption versus number of stations either sending or receiving the chosen flow. Each curve represents a different PER

In the case where the flow is sent from the station to the AP (Fig. 1 and Fig. 2), we can see that power consumption increases linearly with the number of stations. This is because as more and more stations compete for the medium, the collision amount increases as well. And since collisions can not be detected on the half duplex wireless link, the energy required to transmit the whole frame is wasted. In Fig. 1 however, we notice that for five stations and more, the power

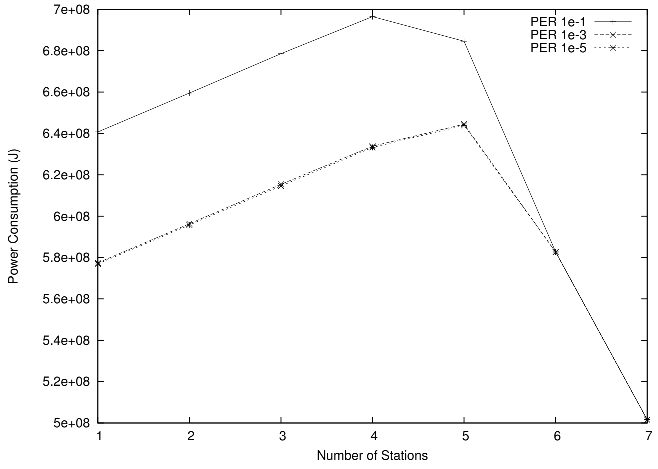

Fig. 1. Station initiated flow, regular mode (power consumption)

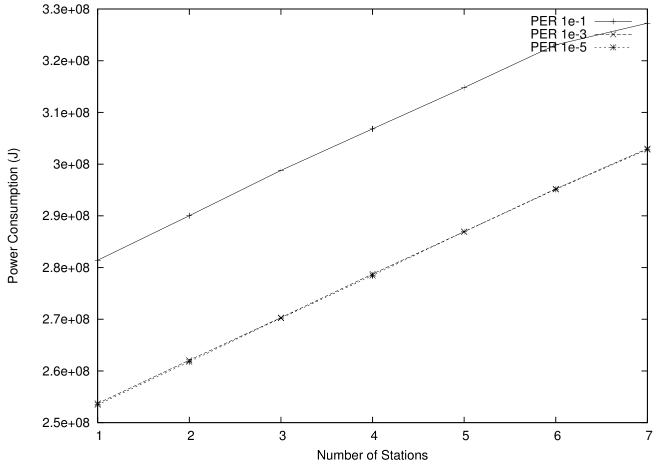

Fig. 2. Station initiated flow, aggregation mode (power consumption)

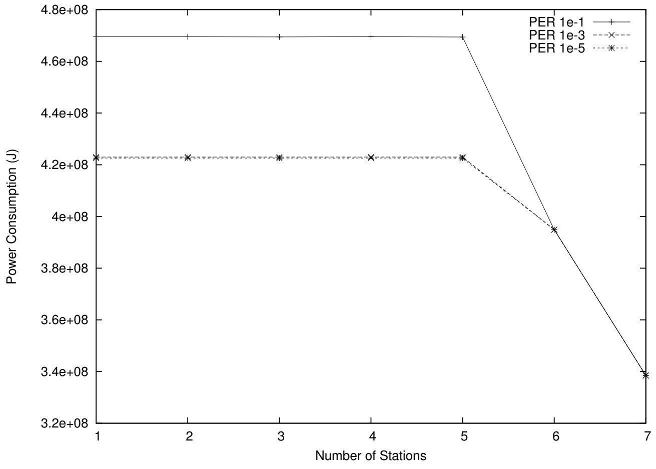

Fig. 3. AP initiated flow, regular mode (power consumption)

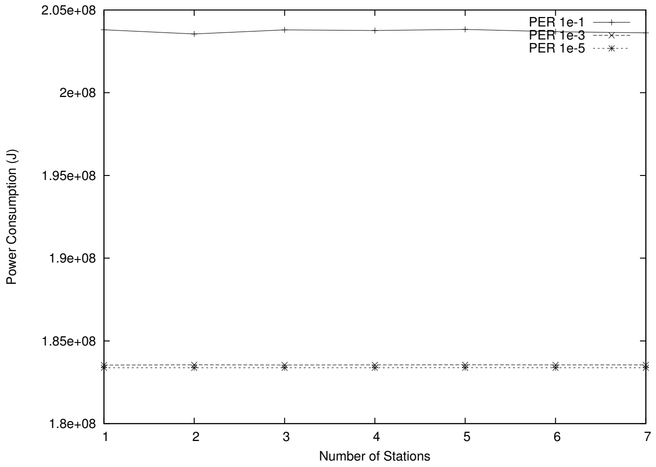

Fig. 4. AP initiated flow, aggregation mode (power consumption)

consumption first increase in a slower way and then drops. We explain this behavior by the fact that for that number of stations, the medium is saturated and many packets are dropped by stations even before trying to send them. In this regard the flow quality is much degraded.

In the other case where the flow originates from the AP, the power consumption from the station standpoint is not affected by the number of stations receiving a flow because the AP is the only transmitter in the cell (see Fig. 3 and Fig. 4). However in Fig. 3 we can see once again that the medium gets saturated. In fact, the saturation occurs for six stations and more because the fairness of the MAC protocol implies that the saturation throughput of one single transmitter (the AP in our case) is lower than the combined throughput of several, but it is always higher than the throughput of any station having to compete with others. The power consumption decreases after saturation point while the number of stations increase because less packets are received by each station. Although the overall number of packets delivered remains the same, each station receives less and thus the power consumption per station drops.

B. Latency Results

For latency results, we can not show average results because each simulation run is unique in essence. Each simulation reports a latency value in microseconds for each frame. An average of such values would be pointless because it would smooth peak latency values caused by retries.

However, we analysed all the results we could get from various scenarios and show here the most representative results for our purpose. The following figures show the plot of latency along the y-axis according to frame sequence number along x-axis. Negative latency values mean that the packet was lost to a frame error, or dropped because the queue of the wireless network interface was full.

We will first have a look at results for one single station transmitting in the cell. This will obviously show the drawbacks of Frame Aggregation because of the additional latency introduced by the mechanism, whereas regular mode performs very well in this environment. For both plots, the loss rate was 10^{-3}. The reader will have noticed that the figure shown in Fig. 6 and Fig. 5 only present an extract from the complete figure because it would be unreadable with the full thirty

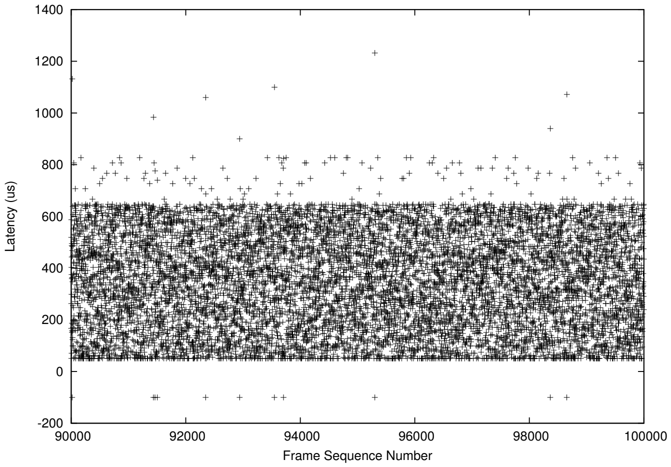

Fig. 5. Station initiated flow, regular mode (latency)

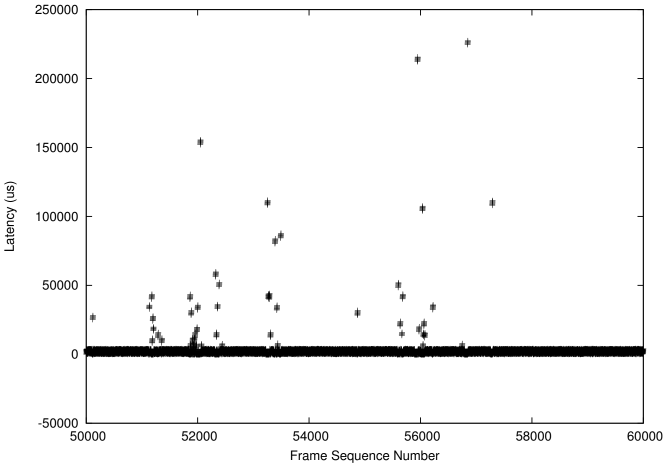

Fig. 7. Station initiated flow, regular mode (latency)

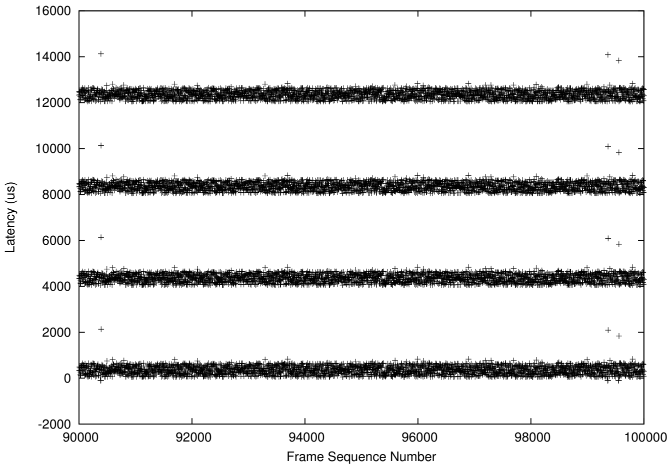

Fig. 6. Station initiated flow, aggregation mode (latency)

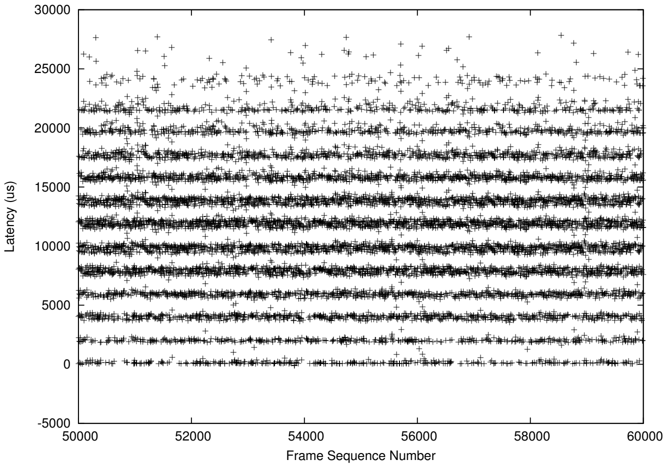

Fig. 8. Station initiated flow, aggregation mode (latency)

minutes of simulated time.

In Fig. 5 the latency stays well below two milliseconds which is a good performance. In comparison, Frame Aggregation has four time worse results because some frames almost reach sixteen milliseconds latency. When taking a closer look at Fig 6 we can notice four stripes where all latencies are concentrated. Each stripe is associated with one of the aggregated frames. The distance between each stripe is the period between each packet in the flow definition. Although Frame Aggregation performs four times worse than regular mode, sixteen milliseconds is still a low latency.

We noticed in previous subsection that using regular mode, the cell saturates as soon as more than four stations transmit at once. We expect the latency to vary accordingly. This can be verified in Fig. 7 where some peaks reach two hundreds milliseconds and more, which is unacceptable for realtime communications. In fact, in the full plot there are some peak values reaching one second latency. This figure plots the latency for each frame sequence number in a saturated state caused by six stations in competition to transmit the flow.

In Fig. 8 we see that for the chosen extract, latencies are spread from almost none up to thirty milliseconds. The simulation environment was the same as in Fig. 7 but saturation does not occur in this case. We can still notice stripes with high density of values caused by aggregation mechanism, but there are more than four because collisions might defer the frame by some interframe spacing time.

C. *Results Analysis and Discussion*

We learned from section IV-A that Frame Aggregation performs better than regular mode especially because it allows more stations to be served regarding medium access. However, we can see from the figures in that section that Frame Aggregation is much more power efficient than regular mode. We

TABLE I

POWER CONSUMPTION RATIO FOR TRANSMITTING STATIONS

Number of STAs	1	2	3	4
Ratio for 10^{-1}PER	2.277	2.274	2.271	2.270
Ratio for 10^{-3}PER	2.276	2.275	2.275	2.274
Ratio for 10^{-5}PER	2.276	2.276	2.274	2.274

TABLE II

POWER CONSUMPTION RATIO FOR RECEIVING STATIONS

Number of STAs	1	2	3	4	5
Ratio for 10^{-1}PER	2.303	2.307	2.303	2.304	2.303
Ratio for 10^{-3}PER	2.304	2.304	2.304	2.304	2.304
Ratio for 10^{-5}PER	2.304	2.304	2.304	2.304	2.304

summarized the ratios in Table I for values corresponding to non saturated states with transmitting stations, and in Table II for receiving stations. And both case, Frame Aggregation is more than twice as efficient as regular mode.

This allows Frame Aggregation to perform better than regular mode even with high PER because the energy expenses of one retransmission for an aggregation superframe will be covered by the savings of the next successful transmission.

However, PER is never a constant in wireless environment. In fact, the real phenomenon is not PER but Bit Error Rate (BER). And in real conditions, BER can occur so that bigger frames can not make it through the wireless medium. Another special case of errors is the burst where many packets are lost in a row. The latter is no more harmful to Frame Aggregation than to regular mode since it would cause nearly as much wasted energy in both cases and the difference would be very slight : if the burst is long enough to corrupt two aggregated superframes, it would surely destroy almost as many regular frames as there were in both aggregated superframes.

From this analysis, we can come up with this adaptive parameter for the Frame Aggregation mechanism : if two consecutive aggregated superframes are lost, we are either facing a high BER or a burst of errors (or both). We then switch to regular mode until the medium becomes clear for long enough, making it sure that the burst is over. If switching back to Frame Aggregation mode our superframes are still undeliverable, we switch back to regular mode and so on.

The performance of this adaptive mechanism will obviously depend on the inhibition period we will choose before switching back to Frame Aggregation mode. Smaller values will make the adaptation more reactive whereas larger values will help to make sure that we do not switch back and forth because of a high frequency BER.

V. CONCLUSION

In this paper, we showed results of extensive simulations performed around IEEE 802.11 cells and our power saving mechanism : Frame Aggregation. The goal of this mechanism is to provide a bidirectional way of saving energy that can handle realtime communications. The simulations results are very encouraging because Frame Aggregation benefits are twofold. It is very power efficient because it consumes much less power than regular operating mode (2:1 ratio) but it allows to maximize the cell capacity too by delaying the occurence of the saturation point where all communications are degraded.

There is a strong requirement [8] in WLANs to provide a power saving mechanism that is latency friendly before they can be used for realtime applications like interactive voice (telephony) or video. In addition to fulfilling these requirements, our proposition retains backwards compatibility with existing hardware, allowing regular stations to operate in a BSS that is compatible with Frame Aggregation (but not to take advantage of Frame Aggregation).

Eventually a key feature of wireless communications is the highly dynamic nature of the environment which requires the MAC protocols to adapt to changing conditions. This is the purpose of the adaptation mechanism we introduced in section IV-C, in which a parameter allows to choose the adaptation rate of the protocol (at the cost of reduced power savings).

REFERENCES

[1] *Part 11 : Wireless Lan Medium Access Control (MAC) And Physical Layer (PHY) Specifications*, Sep 1999.

[2] C. Röhl, H. Woesner, and A. Wolisz, “A short look on power saving mechanisms in the Wireless LAN Standard IEEE 802.11,” *Advances in Wireless Communications*, pp. 219 226, Apr 1998.

[3] J.-C. Chen, K. M. Sivalingam, P. Agrawal, and S. Kishore, “A comparison of MAC protocols for wireless local networks based on battery power consumption,” in *INFOCOM (2)*, Mar. 1998.

[4] J. Lorchat and T. Noel, “Energy Saving in IEEE 802.11 Communications using Frame Aggregation,” in *IEEE Globecom*, 2003.

[5] A. S. M. Simulator, “http://www-r2.u-strasbg.fr/ lorchat/simplemac.”

[6] L. M. Feeney and M. Nilsson, “Investigating the energy consumption of a wireless network interface in an ad hoc networking environment,” in *INFOCOM (2)*, 2001.

[7] B. Burns and J. Ebert, “Power consumption, throughput and packet error measurements of an IEEE 802.11 WLAN Interface,” TKN - TU Berlin, Tech. Rep. TKN-01-007, Aug 2001.

[8] N. Bambos, “Toward power sensitive network architectures in wireless communications : Concepts, issues and design aspects,” *IEEE Personal Communications*, pp. 50 59, Jun 1998.