Approaches To Reduce Energy Consumption of WLAN devices

John-Patrick Wowra
Center for Informatics
University of Goettingen
Email: jpwowra@math.uni-goettingen.de

September 14, 2004

Abstract

Mobile devices power consumption is very high and existing power saving solutions are not perfect. There are different approaches to reduce power consumption of mobile devices.

In this paper two approaches will be presented and evaluated. The first approach reduces power consumption with an application seneitive power management deciding whenever it is necessary to switch into Power Saving Mode. The second approach reduces energy costs of transmissions by computing the optimal strategy to transmit with highest efficiency. Evaluation shows that both attempts work fairly well in reaching the goal of reducing the energy consumption of IEEE 802.11 devices.

1 Introduction

Nowadays mobile connectivity becomes more and more important. Wireless networks provide mobile computers with continuous Internet connectivity. Power management is needed since network interfaces would overly strain the limited battery capacity of a mobile device.

Current wireless network power management often substantially degrades performance and may even increase overall energy usage when used with latency-sensitive applications. Therefore reducing the energy consumption of wireless devices is perhaps the most important issue in the widely deployed IEEE 802.11 Wireless LAN. In this paper two different approaches are presented, both trying to reduce energy consumption of WLAN devices.

The first attempt is called Self Tuning Network Power Management (STPM) and reduces energy consumption by adapting the behaviour of the network device to the access patterns and intent of applications, the characteristics of the network device and the energy usage of the platform.

The second approach is called Minimum Energy transmission strategy (MiSer) trying to minimize the communication energy consumption in 802.11 a/h systems by combining transmission power control with physical layer rate adaptation.

In this paper the two approached are presented [1, 2] and finally compared to each other.

2 Self Tuning Network Power Management (STPM)

As already mentioned the current network power management degrades performance and may increase the overall energy usage. The IEEE 802.11 standard provides at least the following two power modes “CAM” - Continuously aware mode and “PSM” - Power saving mode.

PSM causes an unacceptable 16-32x slow down in the time to list directories stored in Network File Systems in comparison to CAM.

STPM differs substantially from those strategies. It considers the time and energy costs of changing power modes. In fact these time costs can be quite large for current IEEE 802.11 cards, some hundred milliseconds.

Furthermore STPM adapts its behaviour to access patterns and intent of applications. The base power usage of a mobile computer is also considered by STPM. Applications are allowed to disclose hints about their intent in using the network interface with a simple interface provided by STPM. The power management strategy is then adapted to observed network access patterns. By that energy consumption can be decreased while access speed is increased.

Having described the general idea of STPM, in the following part the design principles of the system will be revealed.

2.1 Design Principles

STPM is based upon the following design principles:

2.1.1 Know Application Intent

Only little knowledge about the applicastion intent goes a long way. The most common applications issue file operations sequentially, thus NFS [3] often has only a single remote procedure call in flight. The current IEEE 802.11b power management effectively limits NFS to one remote procedure call per beacon period. Although data rate of NFS is low, its data rate would increase substantially because several remote procedure calls could complete during each beacon period. However PSM does no transition to CAM because it does not detect enough network traffic.

Another strategy would be to switch to CAM whenever an incoming packet is received, but this aggressive strategy works out poorly in other cases.

For example consider a stock ticker application that receives approximately 10 packets per second. When power management is enabled NFS and stock ticker application receive roughly the same amount of data per second. However stock ticker performance will not improve when power management is disabled because it is already receiving at its maximum data rate.

Without knowing the application intent, it is hard to distinguish between these two applications. STPM allows applications to disclose hints about their intent using the wireless network. This allows STPM to enable power management only when appropriate. Further a hint based approach helps STPM decide if the network device can be disabled for periods longer than the beacon period.

If each application discloses hints when it is transferring data and specifies the maximum delay on incoming packet arrivals it is willing to tolerate, the STPM can disable the network device when it is not being used and ensure that the application delay constrains are satisfied.

2.1.2 Be Proactive

Only little knowledge about the application intent goes a long way. The most common applications issue file operations sequentially, thus NFS often has only a single remote procedure call in flight. The current IEEE 802.11b power management effectively limits NFS to one remote procedure call per beacon period. Although data rate of NFS is low, its data rate would increase substantially because several remote procedure calls could complete during each beacon period. However PSM does no transition to CAM because it does not detect enough network traffic. Another strategy would be to switch to CAM whenever an incoming packet is received, but this aggressive strategy works out poorly in other cases [4].

For example consider a stock ticker application that receives approximately 10 packets per second. When power management is enabled NFS and stock ticker application receive roughly the same amount of data per second. However stock ticker performance will not improve when power management is disabled because it is already receiving at its maximum data rate.

Without knowing the application intent, it is hard to distinguish between these two applications. STPM allows applications to disclose hints about their intent using the wireless network. This allows STPM to enable power management only when appropriate. Further a hint based approach helps STPM decide if the network device can be disabled for periods longer than the beacon period.

If each application discloses hints when it is transferring data and specifies the maximum delay on incoming packet arrivals it is willing to tolerate, the STPM can disable the network device when it is not being used and ensure that the application delay constrains are satisfied.

2.1.3 Respect the Critical Path

Latency is often critical when data transfers are driven by an interactive application. The perception threshold beyond which delays become noticeable to human beings is quite small. Typically this threshold is situated between 50 and 200 ms. This means that only a few small transfers in PSM can cause a noticeable delay which may be frustrating users. However there is also a substantial amount of network traffic for which latency is not critical.

For example streaming multimedia applications that buffer data on the client can tolerate delays commensurate with their buffer sizes. To differentiate between these two types of network traffic, STPM enables applications to hint whether a transfer is a foreground transfer in which latency is constrained, or a background transfer that is not time critical. In the first case STPM tries to reduce transfer time and conserve energy, in the second case STPM considers only energy conservation.

2.1.4 Embrace the Performance / Energy Trade-off

If a mobile computers battery is fully charged and the user intents to operate on battery power for only a short time, energy conservation is unnecessary and the user should choose a power management strategy that maximises performance.

However if a mobile computers batteries are nearly exhausted energy conservation is of primary importance.

STPM provides a 'knob' to adjust the relative priority for energy conservation and performance. By that a user is able to distinguish by himself weather performance or energy conservation shall be of priority for STPM.

2.1.5 Adapt to the Operation Environment

To set the correct power management policy, STPM must understand not only the energy characteristics of the network interface but also those of the computer using the interface.

The goal of power management is to extend a mobile computers battery lifetime this means that the energy usage of the entire computer must be minimised, not simply that of the network device.

With incorrect use of network power management, the amount of useful work that a user can accomplish on battery power may decrease. It may also be possible that the correct power management strategy for one device may be inappropriate for another one.

2.2 Characterising Network Power Costs

Wireless network devices differ substantially in the types of power saving modes that are supported and in the power that is used in each mode. Several 802.11b cards have custom adaptive algorithms implemented in firmware.

The power usage of different cards can vary by a factor of two and the transition cost of switching power modes duffer as much as 150 ms. In this approach to educe power consumption a benchmark was created which measures: the base power, when a computer is in idle and no network card attached.

Further it is measured how much power is consumed in each mode (CAM, PSM and others if existent), the transition costs to switch from one mode to another and the average power usage to send and receive 4 MB data in each power mode. This characterisation allows STPM to tune its behaviour to the specific card installed on the system.

2.3 Setting up the Power Management Policy

The STPM algorithm contains 3 rules which prescribe when power modes are to be changed.

1. The system switches when any application specifies a delay tolerance less than the maximum latency of PSM.

2. A transition is also accomplished, when any application discloses that the forthcoming transfer will be large enough, so that the expected cost of performing the transfer in PSM is larger than the expected cost of switching to CAM and then performing the transfer.

3. Finally the system transitions when any application discloses a forthcoming transfer and based on recent access patterns, STPM expects that there will be many short transfers that the cumulative benefit of switching to CAM is greater than the transition cost.

The first case is straightforward and does not need any explanation. When a transfer hint is disclosed, STPM checks for the second case with a cost benefit analysis.

First STPM calculates the total cost of switching to CAM by adding the estimated time and energy necessary to switch to CAM to the transaction costs given by the benchmark results.

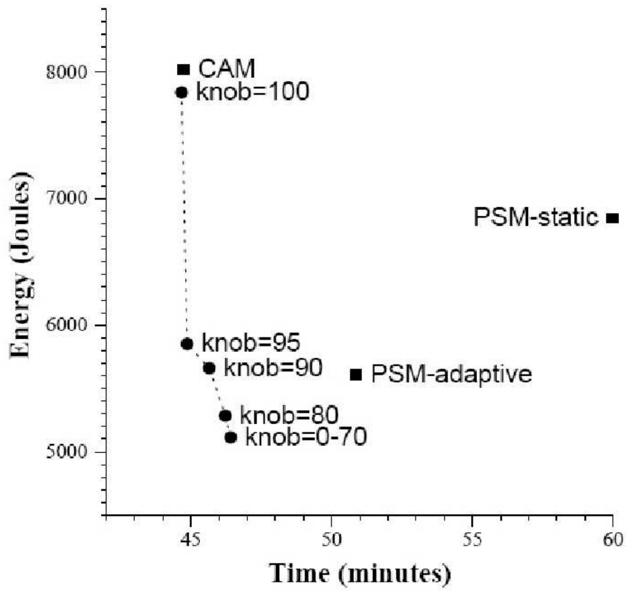

Figure 1: Self-Tuning Wireless Network Power Management

Then the result is compared to the estimated energy to perform the transfer in PSM. Time and energy of a single transfer is insufficient in this case to justify switching to CAM. STPM calculates an empirical probability distribution of transfer hint frequency.

2.4 Evaluation

This approach to reduce energy consumption was investigated in different network intensive application scenarios.

First file access using the Coda distributed file system was measured. Then file access using NFS and playing streaming audio using Xmms. The last scenario is hosting a thin client remote X application. Two of the test scenarios will be presented as examples.

The first figure shows how performance and energy usage vary for the coda [5] scenario depending upon the value of the STPM knob parameter. Each circle represents results using STPM for a different knob value. Knob values of 0 - 70 yield equivalent results. The boxes show the performance and energy usage achieved when native modes of the device are used. The results prove that STPM can reduce power consumption while improving the performance of the device, even when the 'knob' is set to minimum energy consumption the data transfer takes almost five miutes less in comparison to adaptive power saving mode. Static power saving mode is the less energy efficient solution, it takes almost 15 more minutes to fullfill the transfer while

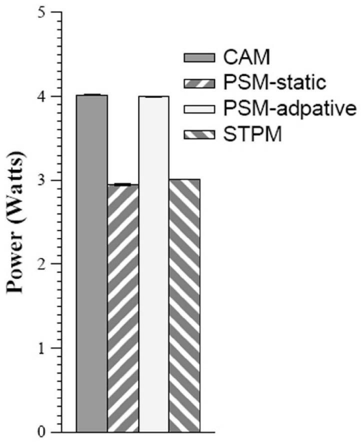

Figure 2: Self-Tuning Wireless Network Power Management

energy consumption is reduced by approximately 12% in comparison to CAM.

The next scenario presented here is play streaming audio using Xmms. The graphs show how the choice of power management policy affects the power used to play streaming audio. Each bar shows the mean of three trials the error bars show the value of the minimum and maximum trial. The results show that STPM works significantely better than adapive power saving mode. Static power saving mode works as well as STPM since STPM will work in the same manner while streaming audio and puffering the date.

2.5 Summary

Wireless network power management can severely degrade performance of latency sensitive applications and increase total energy consumption.

Furthermore it is infeasible to expect a user to tune the power management manually. The results of the evaluation of this approach show that STPM improves performance and energy conservation compared to current power management strategies. Aditionally it does not confuse the user. The power management will work on the background with no need of human attention. It also enables interested users to distinguish whether the priority shall be set to performance or energy conservation.

3 MiSer System Overview

MiSer stands for Minimum energy transmission strategy and is another attempt to reduce power consumption of IEEE 802.11 devices.

A WLAN device can be in different power modes like transmit, receive, idle or doze mode. It consumes the highest power in the transmit mode and very little energy in doze mode. In idle mode the device is required to sense the medium and consumes as much power as it does in receive mode [6].

Several power management policies have been proposed to force an IEEE 802.11 device to enter doze mode adaptively at appropriate moments to save battery power as the attempt to reduce energy consumption presented in the first part of this paper.

An alternative way to conserve energy is to apply transmit power control (TPC) [7, 8, 9] in WLAN systems, which allows a WLAN device to use the minimum required power level in the transmit mode and is complementary to power management policies.

The IEEE 802.11 physical layers provide multiple transmission rates by employing different modulation and channel coding schemes. For example the 802.11b physical layer provides y physical layer rates from 1 to 11 Mbps at the 2.4 GHz band. MiSer is based upon TPC and physical layer rate adaptation (PHY) [10, 11].

The key idea is to compute offline an optimal rate power combination table indexed by data transmission status. Each entry in the table is the optimal rate power combination in sense of maximising energy efficiency under the data transmission status. The data transmission status is characterised by the payload of the data, path loss from transmitter to receiver and frame retry counts. At runtime an energy efficient transmission strategy is determined to minimize the transmission energy costs.

3.1 Table Establishment

In order to establish the rate power combination, a wireless station needs to know the Network configuration that indicates the number contending stations and determines the request to send (RTS) collision probability. Furthermore information about the wireless channel mode is needed, that determines the error performance of physical layer rates. For gaining the required knowledge a simple and effective TPC mechanism is used. Each entry of the table is the optimal rate power combination in sense of maximising the energy efficiency under the corresponding data transmission status.

The data transmission status is characterised by data payload length, path loss from transmitter to receiver and the frame retry counts. The energy efficiency is defined as the ratio of the expected delivered data payload to the expected total energy consumption.

This table is used at runtime to determine the proper physical layer rate and transmit power for each data transmission attempt. When the table is computed offline the transmission error probabilities are negligible because of their smaller frame sizes and robust transmission rates.

It is also assumed that future retransmission will be made with most energy efficient transmission strategies. Frame delivery is only successful, if the request to send transmission succeeds without collision and data transmission is error free or results in correctable errors. Otherwise the station has to re-contend for the medium to re-transmit the frame. The request to send collision probability varies with the network configuration and the data transmission error probability varies with the wireless channel mode.

Since there are only finite choices for the physical layer rate and transmit power the energy efficiency can be calculated for each rate-power combination and the pair that maximises the energy efficiency is the most energy efficient strategy for the data transmission attempt. If at least one of the frame retry limits has been reached, the data frame will be discarded without any further transmission attempt. So the rate power table can be fully computed.

3.2 Runtime Execution

Before running the program the wireless station computes the optimal rate power combination for each set of data payload and length, path loss and frame retry counts. Thus, a rate power combination table is pre-established and ready for runtime use.

At runtime, the wireless station estimates the path loss between itself and the receiver, and then selects the power rate combination for the current data transmission attempt by a simple table lookup. If a request to send / clear to send frame successfully reaches the transmitter and an acknowledgement frame is received correctly within the time limit, the knows that the previous transmission was successful. The wireless station will re-select the rate power combination for the next transmission attempt.

If the frame cannot be successfully delivered after maximum medium reservation attempts or maximum data transmission attempts, the frame will be dropped and the next data unit will be attempted to be transmitted.

The fact that computation burden is shifted offline, simplifies the runtime execution significantly. Therefore embedding MiSer at the MAC layer has little effect upon the performance of higher layer applications, which is a desirable feature for any MAC layer enhancement.

3.3 Path Loss Estimation

To determine the best transmission strategy a mobile station has to estimate the path loss between itself and the receiver. A wireless station is enabled to report its transmit power information in the IEEE 802.11h [12]standard.

The transmit power field simply contains the transmit power (in dBm) used to transmit the frame containing the TPC Report element, while the Link Margin field contains the link margin (in dB) calculated as the ratio of the received signal strength to the minimum desired by the station.

The knowledge of the received signal strength via RSSI (Receive Signal Strength Indicator) is as well as the transmit power via the TCP report element found in the

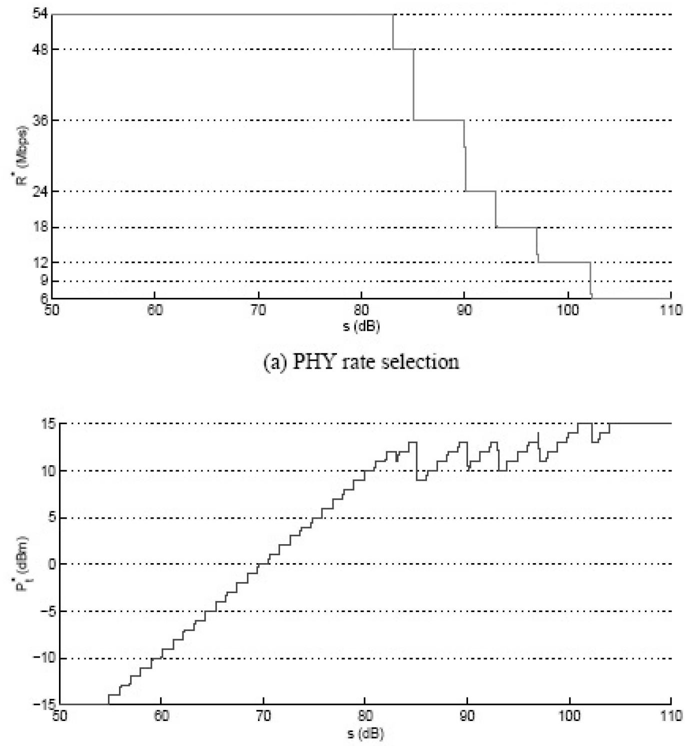

(a) PHY rate selection

(b) transmit power selection

Figure 3: MiSer: An Optimal Low Energy Transmission Strategy for IEEE 802.11 a/h

frame, the wireless station can calculate the path loss from the sending station to itself by performing a simple subtraction.

The RSSI is passed to the MAC by the physical layer and indicates the energy observed at the antenna used to receive the current frame. Since beacon frames are transmitted periodically and frequently, a wireless station is able to update the path loss value in a timely manner.

3.4 Rate Power Combination Table

The optimal combinations of physical layer rate and transmit power, which achieve most energy efficient data communications are shown in the graph under path loss conditions.

As an example, when path loss is 80 dB this figure reads that 54 Mbps, 96 dBm is the most efficient transmission strategy. Two more observations can be made from the graph. When the path loss is large, the lower physical layer rates are preferred as they are more robust and have better error performance. When path loss is small, higher physical layer rates are used to save energy since the duration of a single transmission attempt is shorter.

The second observation shows that a low transmit power does not necessarily save energy. This is because with the same physical layer rate using a lower transmit power may lead to less energy consumption in a single transmission attempt, but the resultant low signal to noise ratio at the receiver side may cause more retransmissions and more energy consumption.

3.5 Evaluation

3.5.1 Star Topologies with Varying Radius

The first testing scenario includes a star topology where eight transmitter stations are evenly spaced on a circle around one common receiver with a radius of 1 to 28 meters. Rate adaptation without transmission power control achieves highest aggregate throughput because its constant use of strong transmit power allows it to choose the highest possible rate to transmit a data frame. On the other hand since rate adaptation has no transmit power control, even within small networks it has to transmit a frame using a higher power than necessary over a small distance, hence consuming more energy. MiSer achieves the highest delivered data per joule because of its adaptive use of the energy efficient combination of high rate and low power when the radius is small and the robust combination low rate and high power when the radius is rather large. The key idea is to select the optimal rate power combination, rather than the physical layer rate or transmit power alone, to minimize the energy consumption. Note that MiSer has the same transmission range as rate adaptation since a transmitter station that supports MiSer can always lower the physical layer rate and / or increase the transmit power to communicate with a far away receiver station.

3.5.2 Random Topologies with 50 Different Scenarios

Testing schemes also include randomly generated network topologies where eight transmitter stations and their receivers were randomly placed within a 40 x 40 meter flat area while all stations are static. During these tests three observations were made. First, MiSer and rate adaptation are significantly better than the single rate transmission power control schemes in each simulated random topology. Second, MiSer achieves comparable aggregate throughput while delivering about 20% more data per unit of energy consumption on average. Third, transmission power control produces near constant aggregate throughput regardless of the network topology. Besides transmission power control has the lowest delivered data per joule in every scenario due to the arbitrary station locations in random topology networks.

3.5.3 Random Vopologies with Varying Mobility

All the testing schemes are relatively insensitive to station mobility. The reason for this insensitivity is that the devices update their path loss conditions to the neighbouring nodes upon each beacon reception, which is every 100ms. Therefore with a maximum speed of 4 m/s the location difference of a wireless station between two path loss updates is 0.2m which has little effect on the path loss conditions and the subsequent rate power selections.

3.5.4 Random topologies with varying data payloads

During these tests payloads simulated are 32, 64, 128, 256, 1024 and 1500 Byte. During the tests both, the aggregate throughput and the delivered data per joule increase with the data payload length for all testing schemes.

MiSer has best energy efficiency performance and the gap between MisSer and rate adaptation becomes larger as the data payload length increases. This is because, with the same physical layer rate, a larger payload results in a longer transmission time, during which MiSer may use low transmit power to save more energy.

3.6 Summary

Based on the observations from the simulation results, the effectiveness of MiSer can be summarized as follows:

MiSer is significantly better than any other scheme that simply adapts the physical layer rate or adjusts the transmit power. Physical layer rate adoption is very effective in saving energy and plays an important role in MiSer.

Applying MiSer does not affect the transmission range while it is insensitive to station mobility and most suitable for data communications with large data payloads.

The effectiveness of MiSer relies on the condition that applying transmit power control on data transmissions will not aggravate the "hidden nodes" problem and the interference in the network. So MiSer exchanges request to send / clear to send frames before each data transmission attempt to deal with the "hidden nodes" problem and transmits the clear to send frames at a stronger power level to ameliorate the interference.

4 Conclusion

Since MiSer is designed as an intelligent transmit power control mechanism used in the transmit mode it is complementary to the power management policies that force a wireless device to enter PSM at appropriate moments to save battery energy. However a simple combination of two may not necessarily result in best energy efficient performance.

For example when the network load is light and the traffic is bursty, a wireless station may want to use physical layer rate adaptation scheme without transmit power control instead of MiSer to finish the frame transmissions as soon as possible and then have the opportunity to enter doze mode earlier to save more energy. This is supported by STPM with the 'knob' to adjust between energy conservation and performance.

In this paper two fairly different approaches were presented trying to solve the problem of energy consumption with wireless devices. STPM can reduce transfer time while saving energy by disabling the device whenever necessary. MiSer reduces the energy consumption by optimiziong the transmit power while not regarding wheater the device is switched on or off.

Thus these two approaches are complementary; a combination of both approaches may result in an even better approach. A device that is only switched on whenever it is needed and transmits at optimal power would significantely improve the quality of wireless LAN devices.

References

[1] M. Anand, E. B. Nightingale, and J. Flinn, "Self-tuning wireless network power management," *MobiCom'03, Sep 14-19, 2003, San Diego, CA USA*, 2003.

[2] D. Qiao, S. Choi, A. Jain, and K. G. Shin, "Miser: An optimal low-energy transmission strategy for ieee 802.11 a/h," *MobiCom'03, Sep 14-19, 2003, San Diego, CA USA*, 2003.

[3] N. W. Group, "Network file system protocol specification," RFC 1094, Tech. Rep., 1989.

[4] R. Krashinsky and H. Balakrishnan, "Minimizing energy for wireless web access with bounded slowdown."

[5] *Disconnected operation in the Coda file system.* ACM Transactions on Computer Systems, 1992.

[6] M.Stemm, P. Gauthier, D. Harada, and R. Katz, "Reducing power consumption of network interfaces in hand-held devices."

[7] S. Agarwal, S. Krishnamurthy, R. Katz, and S. Dao, "Distributed power control in ad-hoc wireless networks."

[8] J.-P. Ebert and A.Wolisz, "Combined tuning of rf power and medium access control for wlans," in *Mobile Networks and Applications.*

[9] J.Gomez, A. Campbell, M. Naghshineh, and C.Bisdikian, "Conserving transmission power in wireless ad hoc networks."

[10] "Ieee 802.11b part 11 wireless lan medium access control and physical layer specifications:high-speed physical layer extension in the 2.4 ghz band," Supplement to the IEEE 802.11 Standard, Tech. Rep., 1999.

[11] "Ieee 802.11a part 11 wireless lan medium access control and physical layer specifications:high-speed physical layer in the 5 ghz band," Supplement to the IEEE 802.11 Standard, Tech. Rep., 1999.

[12] "Ieee 802.11h/d2.2. part 11 wireless lan medium access control and physical layer specifications: Spectrum and transmit power management extensions in the 5 ghz band in europe," Draft Supplement to the IEEE 802.11 Standard-1999 Edition, Tech. Rep., 2002.