

# Wireless LAN Optimization

John-Patrick Wowra

Email: johnpatrickwowra@web.de

Telematics Group  
Institute for Informatics  
University of Göttingen, Germany



# Overview

- Introduction
- Self-Tuning Wireless Network Power Management
  - Design Principles
  - Implementation
  - Evaluation
  - Summary
- MiSer: An Optimal Low-Energy Transmission Strategy
  - Energy Consumption Model & Analysis
  - System Overview
  - Performance Evaluation
  - Summary
- Conclusion

# Introduction

- WLAN problem:
  - Strong power consumption decreases battery lifetime by up to 50%.
  - IEEE 802.11 Power Saving Mode (PSM) is not optimal

# STPM

- Self-Tuning Network Power Management (STPM)
- Motivation
  - Current network power management degrade performance and may increase overall energy usage
  - IEEE 802.11 standard provides at least two modes
    - „CAM“ – Continuously Aware Mode
    - „PSM“ - Power Saving Mode
  - PSM periodically disables the network interface (-> delay)
  - STPM adapts its behavior to access patterns and intent of applications

# Design Principles

- Know application intent
- Be proactive
- Respect the critical path
- Embrace the performance / energy tradeoff
- Adapt to the operation environment

# Know application intent

- Little knowledge about an application helps a lot
  - NFS
    - One RPC per beacon
  - Stock Ticker Application
    - Switch to CAM after receiving a packet not usefull
    - 10 packets per second
    - Performance does not decrease when PSM is used
    - Roughly same amount of data (NFS)
    - Hard to distinguish between these two applications
  - STPN uses hints of the applications about their intent to use the wireless network
  - STPN starts PSM only when appropriate

# Be proactive

- High costs for transition from CAM to PSM
- If data to transfer larger than 4 MB the performance benefit of CAM overweights the transition costs
- NFS is dominated by small transfers
- Network-aware processes like NFS or X disclose start and end of each transfer
- STPM monitors the interarrival time of transfer hints and the number of transfers closely correlated in time (runs)
- STPM supports both, reactive and proactive strategies

# Respect the critical path

- Latency is often critical with interactive apps.
- Perception threshold is cited between 50ms and 200ms
- STPM differentiates between foreground (time-critical) and background traffic
- STPM tries to reduce transfer time and conserve energy

# Embrace the performance / energy tradeoff

- When a device is used for a short time, energy conservation is unnecessary
- When the battery of a mobile device is nearly exhausted, energy conservation is of primary importance
- STPM provides a „knob“ to adjust the priorities

# Adapt to the operating environment

- Energy consumption of the whole computer must be reduced, not only of the network device
- Incorrectly used power management can decrease the amount of useful work a user can accomplish on battery
- The correct power management strategy for one device may be inappropriate for other devices

# Characterizing network power costs

- Power usage of different cards can differ by a factor of two
- Transition costs can differ with different cards by 150 ms
- A Benchmark was created to tune up the module for each network device
- Benchmark measures:
  - » Base power, power when a computer is idle and no network card attached
  - » Power in each mode (CAM, PSM and others if existant)
  - » Transition costs to switch from one mode to another
  - » Average power usage to send and receive 4MB data in each power mode
- Characterisation allows STPM to tune its behavior to the specific card installed on the system

# Setting the power management policy

– STPM transitions from PSM to CAM when:

1. Any application specifies a delay tolerance less than the maximum latency of PSM
2. 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. 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.

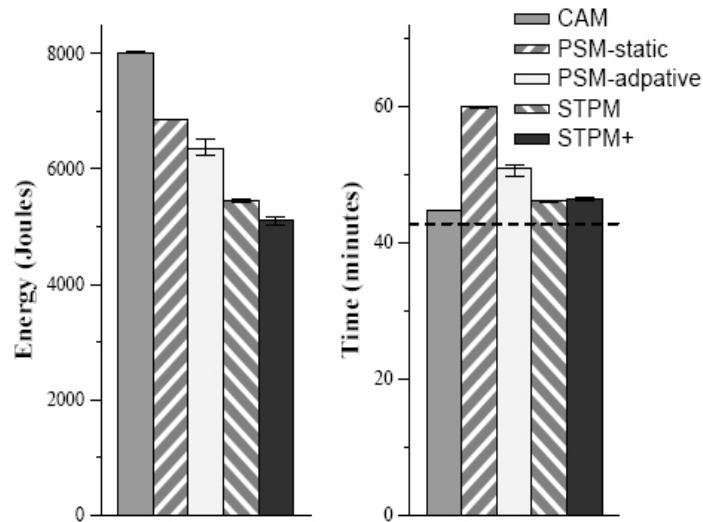
# Setting the power management policy

1. This case is straightforward.
2. When a transfer hint is disclosed, STPM checks for second case with a cost / benefit analysis
  1. STPM calculates the total cost of switching to CAM by adding the estimated time and energy necessary to transfer to CAM to the transaction costs given by the benchmark.
  2. The result is compared to the estimated time and energy to perform the transfer in PSM.
3. Time and Energy of a single transfer is insufficient to justify switching to CAM. STPM calculates an empirical probability distribution of transfer hint frequency.

# Evaluation

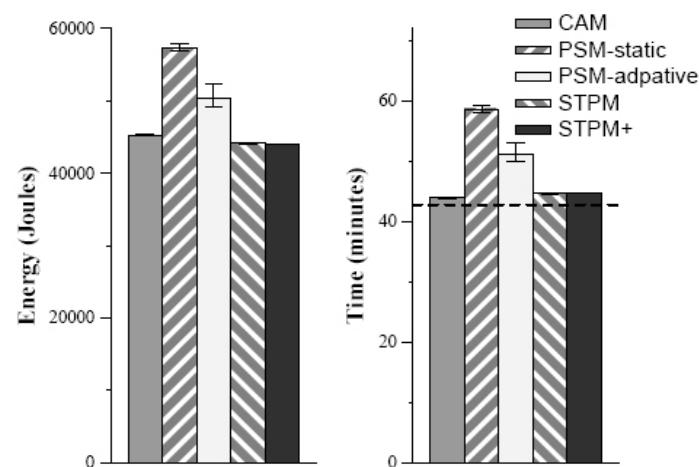
- STPM was investigated using four network-intensive application scenarios
  1. File access using the Coda distributed file system
  2. File access using NFS
  3. Playing streaming audio usinf Xmms
  4. Hosting thin-client remote X application
- PCMCIA card used in tests supports 3 power modes
  - CAM
  - PSM-adaptive
    - » switches between CAM and PSM upon traffic
  - PSM-static

# Coda file system



These graphs show how the choice of power management policy affects the time and energy needed to execute a trace of Coda operations. The dashed line in the rightmost graph shows the amount of think time in the trace. Each bar shows the mean of three trials—the error bars show the value of the minimum and maximum trial.

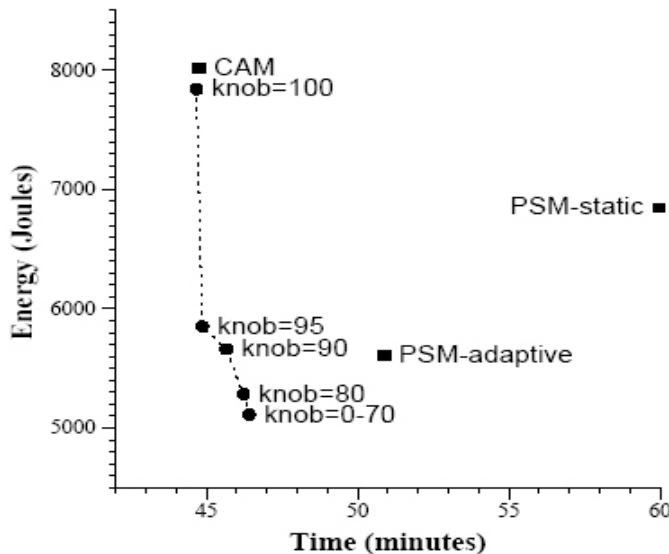
Figure 7: Benefit of STPM for Coda



These graphs show the time and energy costs of executing the Coda trace on an IBM T20 ThinkPad laptop. These results can be compared to those in figure 7 that were collected using an iPAQ client. Each bar shows the mean of three trials—the error bars show the value of the minimum and maximum trial.

Figure 8: Coda results on IBM T20 laptop

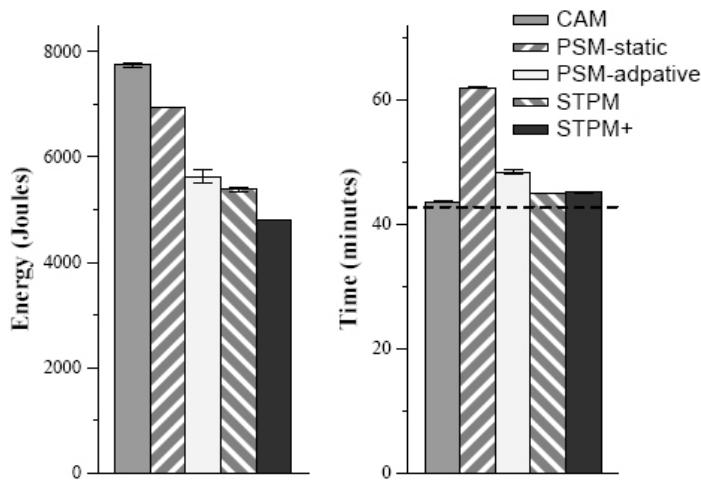
# Coda file system



This figure shows how performance and energy usage vary for the Coda 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 through 70 yield equivalent results. The boxes show the performance and energy usage achieved when the native modes of the Cisco card are used.

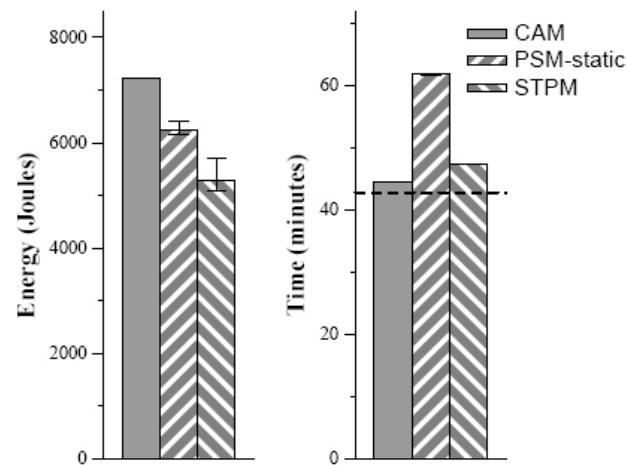
Figure 9: The performance/energy tradeoff

## NFS



These graphs show how the choice of power management policy affects the time and energy needed to execute NFS operations. Each bar shows the mean of three trials—the error bars show the value of the minimum and maximum trial.

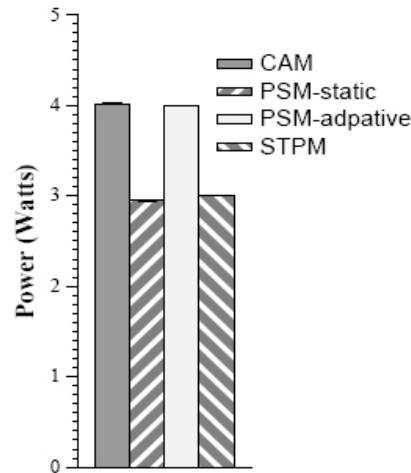
Figure 10: Benefit of STPM for NFS



These graphs show the time and energy needed to execute NFS operations using a Orinoco Silver 802.11b card. These results can be compared to those in figure 10 for the Cisco Aironet 350 card. It is not possible to disable the Orinoco card, nor does it support an adaptive PSM. Each bar shows the mean of three trials—the error bars show the value of the minimum and maximum trial.

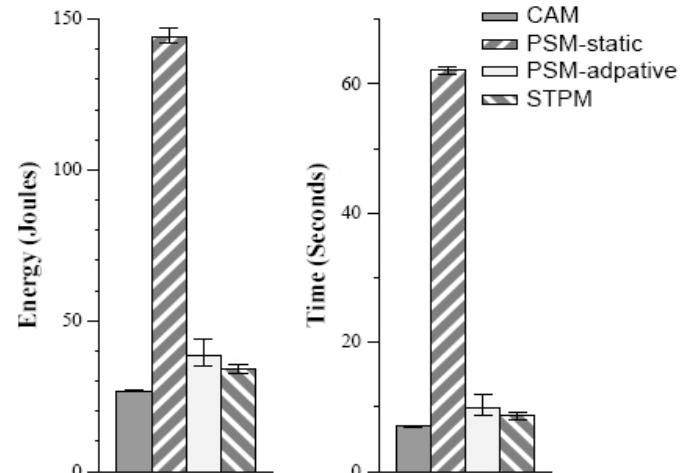
Figure 11: NFS results for Orinoco card

# Streaming Audio (Xmms)



These graphs show how the choice of power management policy affects the power used to play streaming audio on an iPAQ using XMMS. Each bar shows the mean of three trials—the error bars show the value of the minimum and maximum trial.

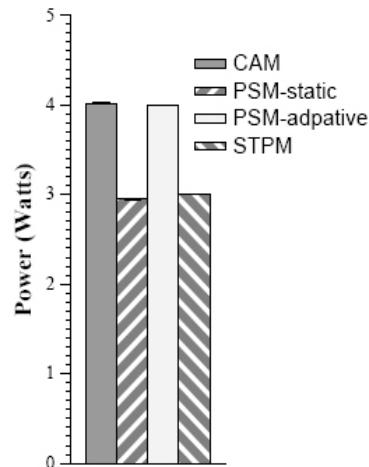
Figure 12: Benefit of STPM for Xmms



These graphs show how the choice of power management policy affects the time and energy needed to run remote X applications. The rightmost graph shows the time to start the Gnumeric application running on the server with the display on the iPAQ client, load a spreadsheet, and close the application—the leftmost graph shows the total energy expended. Each bar shows the mean of three trials—the error bars show the value of the minimum and maximum trial.

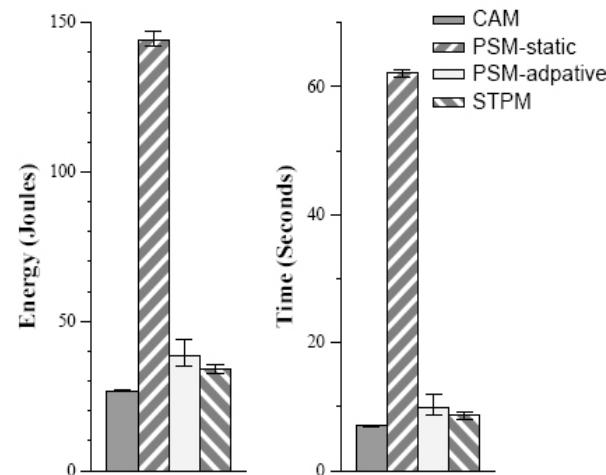
Figure 13: Benefit of STPM for remote X

# Remote X Applications



These graphs show how the choice of power management policy affects the power used to play streaming audio on an iPAQ using XMMS. Each bar shows the mean of three trials—the error bars show the value of the minimum and maximum trial.

Figure 12: Benefit of STPM for Xmms



These graphs show how the choice of power management policy affects the time and energy needed to run remote X applications. The rightmost graph shows the time to start the Gnumeric application running on the server with the display on the iPAQ client, load a spreadsheet, and close the application—the leftmost graph shows the total energy expended. Each bar shows the mean of three trials—the error bars show the value of the minimum and maximum trial.

Figure 13: Benefit of STPM for remote X

# Summary (STPM)

- ***Wireless network power management can severely degrade performance of latency-sensitive applications and increase total energy consumption***
- ***It is infeasible to expect a user to tune the power management manually***
- ***Results show that self-tuning improves performance and energy conservation compared to current power management strategies***

# MiSer

- MiSer is an attempt to reduce power consumption of IEEE 802.11 a/h devices
- MiSer is based upon two technologies
  1. Transmit Power Control (TPC)
  2. Physical Layer Rate adaption (PHY)
- Key idea:
  - Compute offline an optimal rate-power combination table.
  - At runtime an energy efficient transmission strategy is determined

# Energy Consumption Model

- Power Consumption is different for the receive and transmit mode
- High efficiency is achieved at high transmit power levels
- Low efficiency is achieved at low power levels
- Here it is assumed that a WLAN device has the same power consumption in idle mode as in receive mode.

# System Overview

- Basic idea: Wireless stations compute offline a rate-power combination table indexed by data transmission status
- Each entry in the table is the optimal rate-power combination in sense of maximizing energy efficiency under the data transmission status
- Data transmission status is characterised by
  - Data payload length
  - Path loss from transmitter to receiver
  - Frame retry counts

# Runtime Execution

- Following information is needed to compute the table:
  - Network configuration that indicates number of contending stations and determines the RTS (Request To Send) collision detection
  - Wireless channel model that determines the error performance of the physical layer rates
- Path loss estimation
  - Since beacon frames are transmitted periodically and frequently, a wireless station is able to update the path loss value(s) in a timely manner

# Performance Evaluation

- During simulation various network topologies, station mobility patterns and data payload length
- Performances are compared in randomly generated network topologies with 50 scenarios
  - MiSer and RA are significantly better than single rate TBC schemes in both aggregate goodput and delivered data per joule in every simulated random topology.
  - MiSer achieves comparable goodput with RA while delivering about 20% more data per unit of energy consumption
  - TPC / R6 produce near constant aggregate goodput, regardless of network topology

# Performance Evaluation

- Random topologies with varying mobility
  - All the testing schemes are relatively insensitive to station mobility
  - With maximum speed (4 m/s) the location difference of a wireless station between two path loss updates is 0.2 m which has little effect on the path loss conditions and the subsequent rate-power selections
- Random topologies with varying data payloads
  - Simulated date payloads are: 32,64,128,256,512,1024and 1500 Byte
  - MiSer has best energy-efficiency performance
  - Gap between MiSer and becomes bigger as data payload increases

# Summary (MiSer)

- MiSer is clearly better than any other scheme that simply adapts the PHY rate or adjusts the transmit power
- PHY rate adaption is very effective in saving energy and plays an important role in MiSer
- Applying MiSer does not affect transmission range
- MiSer is insensitive to station mobility
- MiSer is most suitable for data communications with large data payloads
- Computation burden is shifted offline
- Embedded MiSer at the MAC layer has little effect on the performance of higher layers

# Conclusion

- Two approaches of different directions to the same problem are hard to compare
- Both ideas are good
- STPM is a more simple approach
- MiSer is trying to solve the problem at the root
- Both systems together may become a very good solution

# The End

Thanks for your patience