

Self-Tuning Wireless Network Power Management

Manish Anand, Edmund B. Nightingale, and Jason Flinn
Department of Electrical Engineering and Computer Science
University of Michigan

ABSTRACT

Current wireless network power management often substantially degrades performance and may even increase overall energy usage when used with latency-sensitive applications. We propose self-tuning power management (STPM) that adapts its behavior to the access patterns and intent of applications, the characteristics of the network interface, and the energy usage of the platform. We have implemented STPM as a Linux kernel module—our results show substantial benefits for distributed file systems, streaming audio, and thin-client applications. Compared to default 802.11b power management, STPM reduces the total energy usage of an iPAQ running the Coda distributed file system by 21% while also reducing interactive file system delay by 80%. Further, STPM adapts to diverse operating conditions: it yields good results on both laptops and handhelds, supports 802.11b network interfaces with substantially different characteristics, and performs well across a range of application network access patterns.

General Terms

Management, Performance

Categories and Subject Descriptors

D.4.4 [Operating Systems]: Communications Management—*network communication*; D.4.8 [Operating Systems]: Performance

Keywords

Power management, self-tuning, 802.11

1. INTRODUCTION

Wireless networks provide mobile computers with continuous Internet connectivity. Yet, power management is needed to ensure that the network interface does not overly tax the limited battery capacity of a mobile device. For example, our measurements show that using a 802.11b network

card without power management can shorten the battery lifetime of a HP iPAQ 3870 handheld by almost 50%.

The popular IEEE 802.11 standard [9] provides a power-saving mode (PSM) that periodically disables the network interface. However, PSM does not adapt to the power characteristics of the network interface and mobile computer, the intent and access patterns of applications, or the needs and expectations of users. While PSM provides excellent energy conservation in some circumstances, it can also substantially degrade interactive application performance and even *increase* the energy needed to perform certain activities. For instance, PSM causes an unacceptable 16-32x slowdown in the time to list directories stored in NFS.

We show that different power management strategies are needed in different circumstances. Rather than take a “one size fits all” approach, we propose self-tuning power management (STPM) that adapts to the characteristics of the network interface, mobile computer, and applications. We have implemented STPM as a Linux kernel module that runs on both handhelds and laptops.

STPM differs substantially from other adaptive strategies such as the PSPCAM mode of the Cisco Aironet 350 card [3] and the bounded slowdown protocol of Krashinsky and Balakrishnan [11]. STPM explicitly considers the time and energy costs of changing power modes. These transition costs can be quite large for current 802.11b cards—several hundred milliseconds in most cases. STPM also explicitly considers the base power usage of the mobile computer. Most importantly, STPM provides a simple interface that allows applications to disclose hints about their intent in using the network interface. STPM then adapts its power management strategy to observed network access patterns.

Our results show that STPM provides significant energy conservation with little performance impact for applications such as distributed file systems, streaming audio, and thin-client remote X displays. For instance, STPM reduces the total energy usage of an iPAQ running the Coda distributed file system by 21% compared to PSM, while also reducing interactive file system delay by 80%. Further, STPM shows benefits across a diverse set of network interfaces and mobile devices.

We begin with a discussion of the limitations of current wireless power management. Section 3 outlines the principles we followed in the design of STPM. In Sections 4 and 5, we describe our implementation and compare its performance and energy conservation to that of other static and adaptive power management strategies. We conclude with a discussion of related and future work.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

MobiCom '03, September 14–19, 2003, San Diego, California, USA.

Copyright 2003 ACM 1-58113-753-2/03/0009 ...\$5.00.

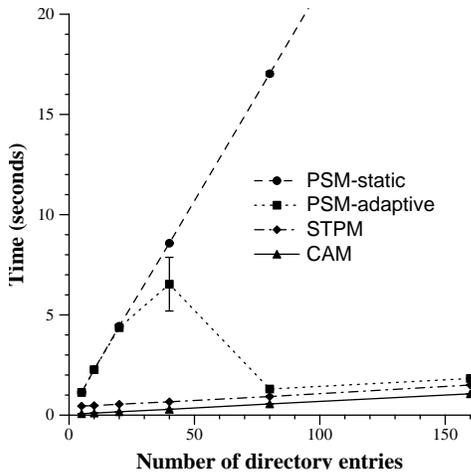


Figure 1: Effect of power management on NFS

2. MOTIVATION

Current 802.11b power management schemes can severely degrade the performance of latency-sensitive applications. For example, Figure 1 shows how power management affects the time to list directories of varying sizes stored in the Network File System (NFS) [17]. These results were generated by typing `ls` on a HP iPAQ 3870 handheld with a Cisco Aironet 350 802.11b card.

The solid line at the bottom labeled “CAM” shows performance in continuously-aware mode (i.e. without power management). The dashed line at the top labeled “PSM-static” shows performance with the default 802.11b power saving mode (PSM). The difference between these two shows that PSM causes an unacceptable 16–32x slowdown for NFS directory listings.

This poor performance is caused by the interaction of NFS remote procedure calls (RPCs) and 802.11b power management. When no packets are waiting for a mobile computer at the wireless access point, 802.11b power management disables the network interface to extend battery lifetime. The access point periodically sends a beacon informing the mobile computer if any packets have arrived—the client interface wakes up to listen to the beacon and goes back to sleep if no data is waiting.

As Figure 2 shows, power management delays each RPC response at the access point until the next beacon. Since NFS does not usually issue concurrent RPCs, only one RPC is issued per beacon period. With a typical 100 ms beacon interval, the expected delay for the first RPC is approximately 50 ms. The delay for subsequent RPCs is almost 100 ms because each RPC request is sent soon after the response to the previous RPC is received. Since NFS makes two RPCs, a `lookup` and a `getattr`, for each file in the directory, the cumulative delay is quite large.

Similar observations have led to the development of adaptive power management strategies that switch between CAM and PSM depending upon traffic load [3, 11]. Ideally, an adaptive strategy can yield good performance by switching to CAM when data is being transferred and by switching back to PSM when it is not. Many cards, including the Cisco Aironet 350 card, support such an adaptive mode.

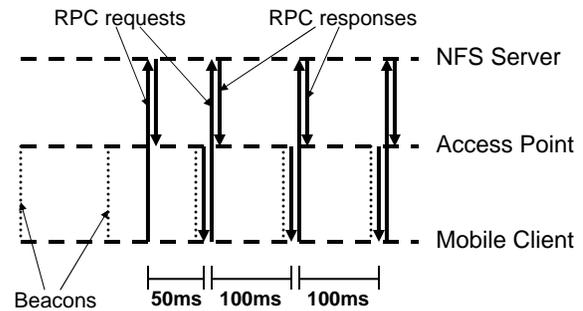


Figure 2: RPC delays due to power management

Our observations of the Cisco card reveal that it switches to CAM when more than one packet is waiting for the mobile computer at the access point, and that it switches back to PSM after approximately 800 ms without receiving a packet.

The dotted line labeled “PSM-adaptive” in Figure 1 shows NFS directory listing performance when the Cisco adaptive mode is employed. For small and medium-sized directories, performance is degraded up to 26x. Because RPCs are issued sequentially and each RPC is typically small, NFS generates insufficient traffic to trigger a switch to CAM. For large directories, the initial read of directory data transmits enough data to trigger the switch—however, even for 160-entry directories, performance still lags CAM by 72% because several RPCs are made in PSM.

Even worse, power management may actually *increase* the energy used to perform interactive tasks. Compared to CAM, the iPAQ uses up to 17 times more energy to list NFS directories with PSM-static and up to 12 times more energy with PSM-adaptive. Although these modes decrease the average power used by the network interface, this decrease is only a portion of the total system power usage. Since energy usage is the integral of power over time, the substantial increase in execution time dominates the small decrease in power. Power management extends battery lifetime but the user accomplishes less work before the battery expires.

Of course, this is a worst-case example. The mobile computer will often be idle for a considerable amount of user think-time; during such idle periods, power management decreases energy usage without performance penalty. Yet, by replaying traces of interactive file system usage, we have found that the cost of using untuned power management during interactive episodes often dominates the benefits realized during idle periods. For many latency-sensitive applications, untuned power management introduces substantial performance and energy penalties. Examples of such applications include other distributed file systems such as AFS [8] and Coda [10], applications that use a remote X server [18] for display, and client-server systems based upon Java RMI [23] and SOAP [1].

To solve this problem, we have built a self-tuning module that adapts to application access patterns, network interface characteristics, and the system on which it is running. The line labeled “STPM” in Figure 10 shows that our module decreases the time needed to list a NFS directory by up to 23x compared to PSM-static and by up to 10x compared to PSM-adaptive. STPM also reduces energy usage by up

to 12x compared to PSM-static and up to 5x compared to PSM-adaptive.

3. DESIGN PRINCIPLES

Self-tuning power management is based upon the following design principles:

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

3.1 Know application intent

A little information about application intent goes a long way. For example, consider why PSM-adaptive works poorly in the previous example. Most common applications issue file operations sequentially; thus, NFS often has only a single RPC in flight. Further, 802.11b power management effectively limits NFS to one RPC per beacon period. Although the data rate of NFS is low, its data rate would increase substantially without power management because several RPCs could complete during each beacon period. However, PSM-adaptive does not transition to CAM because it does not detect enough network traffic.

An alternative strategy would be to switch to CAM whenever an incoming packet is received. However, this aggressive strategy works 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 the stock ticker receive roughly the same amount of data per second. However, the stock ticker performance will not improve when power management is disabled because it is already receiving at its maximum data rate.

Without knowing application intent, it is hard to distinguish these two applications. If an algorithm conservatively refuses to disable power management until a threshold data rate is achieved, it does not disable power management for NFS, leading to poor performance. Alternatively, if it liberally disables power management after the receipt of a few packets, it wastes energy by disabling power management for the stock ticker application.

Our approach is to allow each application to disclose hints about its intent in using the wireless network. This allows STPM to enable power management only when appropriate. Further, a similar hint-based approach helps STPM decide if the network interface can be disabled for periods longer than the beacon period. If each application discloses when it is transferring data and specifies the maximum delay on incoming packet arrivals it is willing to tolerate, then STPM can disable the network interface when it is not being used and ensure that application delay constraints are satisfied.

3.2 Be proactive

If applications such as NFS disclosed hints when each network transfer began and ended, a possible strategy would be to enter CAM whenever at least one transfer is in progress and go back to PSM when no transfers are occurring. This purely *reactive* strategy requires that the transition cost of changing modes be low.

Unfortunately, we have found the transition costs for current 802.11b cards to be quite high. Although the device driver may complete the system call that initiates a mode

transition in only a few tens of milliseconds, packet transmission and reception is delayed for a much longer period of time following each transition. We measured *transition time* for several cards by first initiating a transition to PSM or CAM and then immediately performing a single-packet ping of a nearby server. Transition times ranged from 200 ms to 600 ms—sample results are shown in Figure 6.

A purely reactive strategy increases the time to perform a short RPC because transition time is greater than the latency reduction achieved by performing the RPC in CAM. However, a reactive strategy shows benefits for large requests. For instance, a 4 MB TCP transfer from an iPAQ client with an Orinoco Silver 802.11b card to a nearby server is 16% slower with power management enabled. These results, which confirm previously reported TCP throughput results [11], indicate that there is a break-even transfer size—for transfers larger than this size, the performance benefit of CAM outweighs the transition cost. The particular break-even size is dependent upon the data rate supported by the 802.11b card in each mode and the card's transition costs. STPM determines the break-even point for each card and switches to CAM when an application discloses that a forthcoming transfer will exceed the break-even size.

However, applications such as NFS are dominated by small transfers. For such applications, a proactive strategy is needed to amortize transition costs across multiple transfers. When a proactive strategy determines that a large number of transfers will soon occur, it switches to CAM, and then switches back to PSM after the last transfer. In the NFS example, the cumulative reduction in latency across all RPCs far exceeds the transition cost of changing power modes.

Clearly, the difficulty in implementing a proactive strategy is that it requires foreknowledge of the number of transfers that will occur in the near future. Applications like NFS and the X server do not have this information because they receive each application request sequentially. One possible approach would be for applications like `ls` and `make` to provide hints of future network activity. We rejected this approach because it requires modification of programs that are not normally network-aware. Further, since such applications do not usually care which type of file system they are using and do not know which blocks a file system may have cached locally, they do not know which requests will cause network activity.

Our approach is to have network-aware processes like NFS and the X server simply disclose the start and end of each transfer. STPM monitors the interarrival time of transfer hints, as well as the number of transfers that are closely correlated in time—we refer to such clusters of transfers as *runs*. Using an empirically-collected distribution of run lengths, STPM calculates the expected number of transfers in the current run given the number that it has already been seen. It then performs a cost-benefit analysis to determine if it should switch to CAM. For example, STPM might decide to switch to CAM after three short transfers occur close together, and also to switch back to PSM when 300 ms pass without a further transfer. The details of how this decision is made are explained in Section 4.3.

It is important to note that STPM supports both reactive and proactive strategies. For example, consider a hypothetical network interface that has negligible transition costs. Since the break-even transfer size is effectively zero, STPM would switch to CAM at the start of each transfer and return

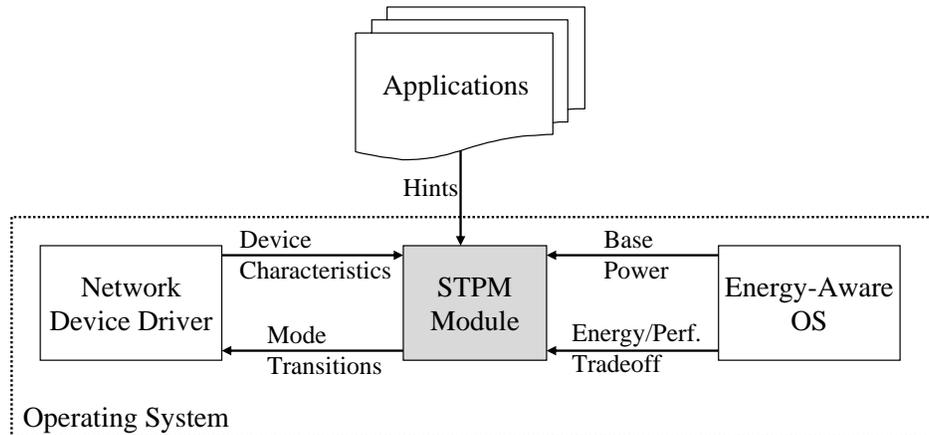


Figure 3: STPM software architecture

to PSM when the transfer completes. Although STPM’s proactive mechanisms would be unused, STPM would benefit fully from the low transition costs.

3.3 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 users is quite small—typically it is cited as being between 50 ms and 200 ms [5, 13]. This means that only a few small transfers in PSM can cause a noticeable delay, and that the cumulative delay for operations such as NFS directory listings may certainly prove frustrating to the user. Thus, it is critical to disable power management when the network is being used by an interactive application.

However, there is also a substantial amount of network traffic for which latency is not critical. For example, the Coda distributed file system prefetches file data from servers to improve performance and guard against disconnection [10]. Coda also writes file modifications back to the server asynchronously [16]. For both prefetching and asynchronous writes, latency is not a critical constraint since a human user is not waiting for the transfer to complete. Similarly, 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 a constraint, or a *background* transfer that is not time-critical. In the former case, STPM tries to both reduce transfer time and conserve energy—in the latter case, STPM considers only energy conservation.

3.4 Embrace the performance/energy tradeoff

Disabling the 802.11 wireless interface reduces power consumption but increases the latency of packet delivery, creating an inherent tradeoff between performance and energy conservation. While this tradeoff seems unavoidable, it is important to evaluate it in the context of a mobile user’s activity. If the mobile computer has a fully-charged battery and the user intends to operate on battery power for only

a short time, then energy conservation is unnecessary, and the user should choose a power management strategy that maximizes performance. However, if the mobile computer’s battery is nearly exhausted, then energy conservation is of primary importance.

A power management policy that statically balances these two competing goals cannot be correct in both contexts. Instead, a tunable strategy is needed. STPM provides a “knob” that can be adjusted to reflect different relative priorities for energy conservation and performance. Users can set the knob to maximum performance when they intend to operate on battery power for only a few minutes—STPM responds by keeping the wireless interface continuously active to minimize latency. We also envision that STPM’s knob might be set by higher-level energy-aware OS components such as Ecosystem [25] or Odyssey [6].

3.5 Adapt to the operating 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 computer’s battery lifetime—this means that the energy usage of the entire computer must be minimized, not simply that of the network interface.

To illustrate the difference, consider a hypothetical power management mode that reduces network power usage by 50% from 2 Watts to 1 Watt, but delays interactive activities by 10%. If this mode were employed by a handheld with a base power usage of 2 Watts, the total power of the mobile computer and interface would be reduced by 25% from 4 Watts to 3 Watts. Since interactive activities now take 10% longer to complete, the total energy used for each activity would be reduced by a still respectable 17.5%. However, if this mode were employed by a laptop with a base power usage of 15 Watts, total power would only be reduced by 5.9%. Further, the total energy used to perform interactive activities would actually increase by 3.5%.

This example illustrates two points. First, when used incorrectly, network power management can decrease the amount of useful work that a user can accomplish on battery

<code>TransferHintBegin</code>	(IN <code>fg_flag</code> , IN <code>xmit_size</code> , IN <code>recv_size</code> , OUT <code>hint_id</code>);
<code>ListenHintBegin</code>	(IN <code>max_delay</code> , OUT <code>hint_id</code>);
<code>HintEnd</code>	(IN <code>hint_id</code>);
<code>SetKnob</code>	(IN <code>knob_value</code>);
<code>SetBasePower</code>	(IN <code>base_power</code>)

Figure 4: Wireless power management API

power. Second, the correct power management strategy for one device may be inappropriate for other devices. STPM avoids these pitfalls by explicitly considering the base power usage of the mobile computer.

4. IMPLEMENTATION

As shown in Figure 3, we have implemented self-tuning wireless network power management as a Linux loadable kernel module. Applications link with a user-level library and disclose hints to the STPM module about their intent and activities. The library implements the STPM API described in Section 4.1 by first opening a Linux pseudo-device and then making `ioctl`s on the device whenever an application calls a STPM function. This implementation enables the STPM module to detect when applications terminate and cancel any outstanding hints that the application made before exiting. Further, the kernel implementation improves performance by minimizing user-kernel boundary crossings.

The STPM module has two further inputs. First, the base power and current tradeoff between energy conservation and performance may be set by an energy-aware operating system [6, 25] or by a user-level configuration tool. Second, for each wireless network card, the STPM module loads a device-specific characterization of power usage and transition costs. Section 4.2 describes how we derive these characterizations by executing a benchmark suite. The characterization is instantiated as a code component that is loaded by the STPM module in much the same way that the operating system loads the device driver for the wireless network card. Currently, the creation of each card-specific component is a manual task, but we plan to have our benchmark suite automatically create such components in the future.

Section 4.3 details the STPM algorithm that decides when to transition the power mode of the wireless network card. This algorithm is the simplest we could envision that meets the design principles of Section 3.

4.1 API

Figure 4 shows the STPM API. Applications use the first three functions to disclose hints about wireless network usage. Before each data transfer, an application calls `TransferHintBegin` and specifies whether the forthcoming transfer represents background or foreground activity. The application may optionally specify the expected amount of data to be sent and received. If the application leaves these values unspecified, the forthcoming transfer is assumed to be small. `TransferHintBegin` returns a unique identifier that the application later passes to `HintEnd` to specify that the transfer has completed. When an application terminates, the STPM module ends any outstanding hints for which the application has not called `HintEnd`.

Processes that are listening for incoming packets may call `ListenHintBegin` to specify the maximum delay due to power

Card Mode	Cisco Aironet 350	Orinoco Silver
$P_{disabled}$	0.24 W	N/A
P_{PSM_idle}	0.39 W	0.19 W
P_{PSM_recv}	1.42 W	2.22 W
P_{PSM_send}	2.48 W	2.70 W
P_{CAM_idle}	1.41 W	1.21 W
P_{CAM_recv}	2.61 W	2.25 W
P_{CAM_send}	3.69 W	2.67 W

This figure shows the power usage of two 802.11 cards in several power modes. From top to bottom, the modes are: card disabled; card in PSM when idle, receiving, and sending data; and card in CAM when idle, receiving, and sending data. The standard deviation of all measurements is within 0.05 W.

Figure 5: Power usage of two 802.11 cards

management that they are willing to tolerate for incoming packets. This function also returns a unique hint identifier that may be passed to `HintEnd` to terminate the listen hint.

The `SetKnob` function adjusts the relative importance of performance and energy conservation. The value of the knob ranges from 0 to 100, where 0 represents maximum energy conservation and 100 represents maximum performance. The `SetBasePower` function specifies the base power usage of the mobile computer.

4.2 Characterizing network power costs

Wireless network interfaces 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—examples include the Cisco Aironet 350 and Intel PRO/Wireless 2011B cards. In addition, we have found that the power usage of different cards can vary by a factor of two and that the transition cost of switching power modes differs by as much as 150 ms.

Our approach to handling this variability is to create a benchmark suite that can characterize the power usage of each new network interface card that we encounter. While running the benchmark suite, we use a digital multimeter to sample the current drawn by a mobile computer. We remove the batteries from the computer and sample current on the input power line. Since voltage is almost constant when power is drawn from the external power supply, current measurements alone are sufficient to determine power usage. Current samples are collected approximately 50 times per second. Many of our benchmarks require communication with a remote computer. For this purpose, we run a server program on a nearby computer that is LAN-connected with the wireless access point. The server program sends and receives bytes when requested.

The benchmark suite first measures the *base power* of the mobile computer—this is the power used when the machine is idle and the network card is not inserted. For each benchmark activity, we subtract the base power from the measured power usage to derive the additional power consumed by the activity.

Next, the benchmark suite measures the power used by the card in each mode. Every card that we have encountered supports at least CAM and PSM. Some support further adaptive modes that switch between CAM and PSM

	Cisco Aironet 350		Orinoco Silver	
	Time	Energy	Time	Energy
PSM to CAM	0.40 s	0.51 J	0.23 s	0.24 J
CAM to PSM	0.41 s	0.53 J	0.26 s	0.31 J
Disable	0.00 s	0.00 J	N/A	N/A
Enable	0.39 s	0.51 J	N/A	N/A

This figure shows the time and energy to transition between CAM and PSM for each card. For the Cisco card, the time and energy to disable and re-enable the card are also shown. The standard deviation of all measurements is within 5%.

Figure 6: Cost of mode transitions

depending upon recent incoming traffic load. It may also be possible to disable the interface without turning the card completely off. For example, when the Cisco Aironet 350 interface is disabled, card power usage is reduced by 150 mW. This saves less power than turning the card off completely but requires less time and energy for transitions.

The benchmark suite characterizes the cost of transitions between each supported power mode. It first initiates a mode transition and then exchanges a one byte ping with the server. The reported transition time is measured from the start of the mode change operation to the receipt of the ping response. The energy cost of the transition is the energy consumed by the network card during this interval (the cost of sending and receiving a byte is not significant within the granularity of our measurements).

The benchmark suite completes the characterization by measuring the power used to send and receive data in each power mode. It first sends 4 MB of data to the server using TCP and measures the average power usage. It then measures the average power used when the server retransmits the buffer to the client. Additionally, it measures the time to perform each transfer, which gives us an estimate of the maximum data rate that can be achieved in each power mode. Figures 5 and 6 show the results of running our benchmark suite for the Cisco Aironet 350 and Orinoco Silver PCMCIA cards.

Card characterization allows STPM to tune its behavior to the specific type of network card being employed. This modular approach means that we do not need to modify the STPM algorithm for each new type of wireless card. The benchmark suite takes approximately two hours to execute—however, characterization need only be done once for each model of wireless card. Ideally, a card manufacturer could run the benchmark suite and provide the characterization as part of the device driver.

4.3 Setting the power management policy

In this section, we describe the STPM algorithm. For the purpose of discussion, we first assume that the network card supports only two power modes: CAM and PSM. For such cards, STPM must decide when to transition to CAM and when to transition back to PSM. The next two sections describe how STPM makes these decisions. Section 4.3.3 then discusses how we generalize STPM to support cards with more power modes.

4.3.1 Transition to CAM

STPM transitions from PSM to CAM when:

- any application specifies a delay tolerance less than the maximum latency of PSM.
- any application discloses that the forthcoming transfer will be large enough such 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.
- any application discloses a forthcoming transfer and, based on recent access patterns, STPM expects that there will be enough subsequent short transfers that the cumulative benefit of switching to CAM for the forthcoming transfers will be greater than the transition cost.

The first case is straightforward. The maximum delay of PSM is equal to the beacon interval, typically 100 ms. If a listen hint is specified that is less than the beacon interval, STPM switches from PSM to CAM.

Whenever a new transfer hint is disclosed, STPM checks for the second case. STPM performs a cost/benefit analysis by estimating the time and energy to perform the forthcoming transfer in both PSM and CAM. Through the transfer hint, the application has disclosed the expected number of bytes to send, B_s , and receive, B_r . The expected time, T , to perform the transfer is:

$$T = L + B_s/DR_s + B_r/DR_r \quad (1)$$

where the expected data rate for sending, DR_s , and receiving, DR_r , data in each mode is given by the card-specific characterization—this assumes that the wireless link is the bandwidth bottleneck in the transfer. The expected latency, L , is measured directly for CAM by observing the shortest time to complete a transfer over the recent past. For PSM, the expected latency is the CAM latency plus half the beacon interval.

The expected energy, E , to perform the transfer is calculated using the specified base power of the mobile computer, P_b , and the measured power usage of the card when idle, P_i , sending data, P_s , and receiving data, P_r . For each mode:

$$E = L \times (P_i + P_b) + B_s/DR_s \times (P_s + P_b) + B_r/DR_r \times (P_r + P_b) \quad (2)$$

Since Equation 2 includes base power, it estimates the total energy used by the entire mobile computer, not just the network interface.

STPM calculates the total cost of switching to CAM by adding the estimated time and energy to perform the transfer in CAM to the transition costs given by the card-specific characterization. It compares the results to the estimated time and energy to perform the transfer in PSM. STPM transitions the card when it predicts that doing so will both save energy and improve performance. However, sometimes STPM estimates that a transition will improve either performance or energy conservation while hurting the other goal. In such cases, it uses the value of the knob that specifies the relative tradeoff between performance and energy conservation to decide whether to transition the card. Given n strategies, each of which has estimated time and energy costs, T_n and E_n , STPM first calculates the mean time, \bar{T} , and energy, \bar{E} , used by all strategies being compared. It

then calculates a relative cost for each strategy, C_n as:

$$C_n = (T_n/\bar{T}) \times knob + (E_n/\bar{E}) \times (100 - knob) \quad (3)$$

For foreground hints, *knob* is the specified relative tradeoff between performance and energy conservation, ranging in value between 0 and 100. The intuition is that the time and energy values are first normalized by dividing them by the mean of the time or energy values being compared, then the knob is used to assign a relative weight for determining the final cost of the mode. For background hints, the knob is set to 0—since the transfer is latency-insensitive, only energy usage is considered.

In the third case, the time and energy of a single transfer is insufficient to justify switching to CAM. However, several subsequent transfers are expected and the time and energy saved over all transfers is expected to exceed the transition cost. To make this determination, STPM estimates the likelihood that subsequent transfers will occur in the near future.

STPM generates an empirical probability distribution of transfer hint frequency by observing the arrival of transfer hints. It maintains a histogram of the number of foreground transfer hints that occur closely correlated together in time. We refer to each such group of correlated hints as a *run*. A run begins when the first transfer hint is issued and ends when 150 ms pass with no foreground transfer being in progress. We chose 150 ms to differentiate the communication patterns of programs such as Web browsers that are driven by a human user from those of programs such as NFS that issue multiple sequential requests without human intervention. Our goal is to have each run correspond to a single interactive activity such as a NFS directory listing.

The *run length* is the number of foreground transfers issued during a run. STPM maintains a 1024 bucket histogram of observed run lengths—if a run exceeds 1024 transfers, STPM records it as having length 1024. The histogram data is periodically read from a `/proc` file system interface and saved to disk by a user-level daemon. When the module is first loaded, the daemon provides the initial histogram values from the saved data. In this fashion, STPM maintains data across reboots.

Currently, histogram entries are persistent, but we plan to investigate how histogram values can be aged to allow the module to adapt quicker to changes in network access patterns. We also plan to investigate the benefits of maintaining per-application histograms. The current approach of maintaining a single histogram for the client allows STPM to aggregate the access patterns of concurrently executing applications; yet, we feel that per-application histograms might allow STPM to adapt better to changing workloads.

We use the histogram to calculate the expected cost of switching to CAM prior to the n th transfer in each run. The expected time, T_n , to execute a run if a switch to CAM is done before the n th transfer is:

$$T_n = \sum_{i=1}^{n-1} L_{PSM} \times P(r \geq i) + \sum_{i=n}^{1024} L_{CAM} \times P(r \geq i) + T_{TC} \times P(r \geq n) \quad (4)$$

$P(r \geq x)$ is the probability that a run will equal or exceed

length x —STPM derives this value from the run length histogram. T_{TC} is the time to switch to CAM. The intuition behind this equation is that STPM adds the total expected latency for transfers performed in PSM, the expected latency for transfers in CAM, and the expected transition time to derive the total expected time to execute the run of transfers. Similarly, STPM calculates the expected energy usage, E_n , as:

$$E_n = \sum_{i=1}^{n-1} L_{PSM} \times (P_{PSM_idle} + P_b) \times P(r \geq i) + \sum_{i=n}^{1024} L_{CAM} \times (P_{CAM_idle} + P_b) \times P(r \geq i) + (E_{TC} + (T_{TC} \times P_b)) \times P(r \geq n) \quad (5)$$

Using Equations 4 and 5, STPM calculates the expected time and energy if it switches to CAM after an application issues the n th foreground transfer hint in a run. Additionally, STPM calculates the expected time and energy if it never switches to CAM. It uses Equation 3 to compare the different policies and chooses the one that has the minimum cost. Since this calculation is relatively time-consuming (2 ms) and the input data is slow to change, STPM performs this calculation every 10 minutes.

4.3.2 Transfer to PSM

STPM transitions from CAM to PSM when no transfers are in progress, no application has specified a delay tolerance less than the maximum latency of PSM, and it estimates that the network interface will be idle long enough to overcome the transition costs of switching modes. To aid in this decision, STPM maintains a 1024 bucket histogram of the length of the interval between each run. Each bucket corresponds to a 100 ms period; extremely long intervals are recorded as having the maximum value of 1024 seconds. The interval histogram is maintained in an identical fashion to the run length histogram described in the previous section.

STPM transitions to PSM when it believes that the expected benefit of reduced power usage during the forthcoming idle period most exceeds the performance and energy cost of beginning the next run in PSM instead of CAM.

The expected time and energy to perform the next run starting in PSM, T_{init_PSM} and E_{init_PSM} , have already been calculated—these values are the expected time and energy of the policy chosen in Section 4.3.1. The expected time and energy to perform the next run starting in CAM, T_{init_CAM} and E_{init_CAM} , are:

$$T_{init_CAM} = \sum_{i=1}^{1024} L_{CAM} \times P(r \geq i) \quad (6)$$

$$E_{init_CAM} = \sum_{i=1}^{1024} L_{CAM} \times (P_{CAM_idle} + P_b) \times P(r \geq i) \quad (7)$$

These equations assume small data transfers. Therefore, transfer time is dominated by latency, L_{CAM} , and card power usage is assumed to be close to its measured idle usage, P_{CAM_idle} . STPM next calculates the expected time and energy costs of switching to PSM $n/10$ seconds after the end of the previous run—these values correspond to the

100 ms histogram buckets. The expected time is:

$$T_n = \sum_{i=1}^{1024} 0.1 \times P(l \geq i) + T_{init_PSM} \times P(l \geq n) + T_{init_CAM} \times P(l < n) \quad (8)$$

STPM calculates $P(l \geq n)$, the probability that the interval length will be greater than $n/10$ seconds, using the interval histogram. The first line of the previous equation is the expected length of the current interval, and the second line is the expected time to perform the succeeding run. The energy cost of switching to PSM $n/10$ seconds after the end of the previous run is calculated similarly:

$$E_n = \sum_{i=1}^{n-1} 0.1 \times (P_{CAM_idle} + P_b) \times P(l \geq i) + \sum_{i=n}^{1024} 0.1 \times (P_{PSM_idle} + P_b) \times P(l \geq i) + E_{init_PSM} \times P(l \geq n) + E_{init_CAM} \times P(l < n) \quad (9)$$

STPM calculates the expected time and energy costs of switching to PSM at each 100 ms interval, as well as the expected cost of remaining in CAM. It then chooses the policy that minimizes Equation 3. This calculation is also performed once every ten minutes, at the same time that STPM decides when to switch to CAM.

4.3.3 Generalizing the model

The previous two sections have shown how STPM creates a transition policy for a card that supports only PSM and CAM. For cards that support more than two modes, there are many more possible policies from which STPM could choose. For instance, after a period of time with no transfers, STPM could transition the Cisco card from CAM to PSM, it could disable the card, or it could decide to switch to PSM and later disable the card if no further transfers occur for another period of time.

Since the number of possible strategies grows exponentially with the number of modes, we employ a heuristic to limit the search space. STPM first calculates the lowest cost policy that transitions to each mode. It then calculates the lowest cost hybrid policies that switch to one mode at the time calculated for the single-switch policy, then make a further transition at some later time. For example, STPM might decide that the lowest cost single-transition policy is to switch to PSM after 300 ms. It would then consider hybrid policies that switch to PSM after 300 ms and disable the card after some further period of time. Since most cards support only a few power modes, this strategy is computationally feasible.

5. EVALUATION

How much does self-tuning power management improve application performance and extend mobile computer battery lifetime?

To answer this question, we modified several network-intensive applications to disclose hints and measured application performance and energy usage when our self-tuning power management algorithm was used. We compared these results to those achieved using the static and adaptive power management algorithms natively supported by the 802.11b cards in Figure 5.

5.1 Methodology

The primary client platform for our evaluation is a HP iPAQ 3870 running the Linux 2.4.18-rmk3 kernel. This handheld computer has a 206 MHz StrongArm processor, 64 MB of DRAM, and 32 MB of flash memory. The measured base power of the iPAQ (when idle with no network card inserted) is 1.44 Watts. Unless otherwise noted, the client uses the Cisco 802.11b card described in Figure 5 to communicate with a Cisco Aironet 350 wireless access point with 100 ms beacon interval. The server in our experiments is a Dell Precision 350 with 3.06 GHz processor and 1 GB DRAM running the Linux 2.4.18-19.8.0 kernel—the server and access point are connected with a 100 Mb/s switch.

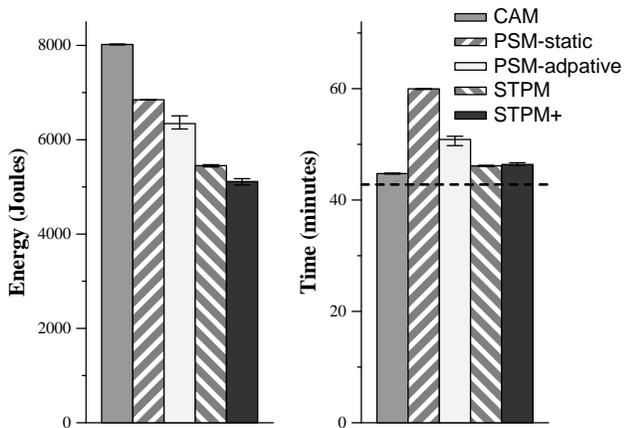
The Cisco PCMCIA card supports three power modes: CAM, where no power management is used; PSM-static, where the card’s receiver is periodically disabled to save power; and PSM-adaptive, which switches between PSM-static and CAM depending upon the incoming traffic load. Since PSM-adaptive is implemented in the card firmware, it can change power modes faster than our STPM module—our measurements indicate that PSM-adaptive almost always transitions in less than 100 ms. In contrast, because STPM is implemented as a kernel module, its transition costs are approximately 400 ms.

We investigated the benefits of self-tuning power management for four network-intensive application scenarios: file access using the Coda distributed file system, file access using NFS, playing streaming audio using Xmms, and hosting thin-client remote X applications. We first ran each scenario using the power management methods natively supported by the 802.11b card. We then executed each scenario using STPM.

Unless otherwise noted, the tuning knob for STPM was set at 50 to equally weight performance and energy conservation. We also warmed the STPM prediction algorithm by specifying an initial probability distribution for network accesses. The alternative approach, starting with no access history, leads to a short period of initial volatility that makes experiments less repeatable. Further, the warming approach better reflects STPM steady-state performance.

To generate the initial access distribution, we replayed a trace of distributed file system accesses using Mummert’s DFSTrace tool [16]. This tool re-executes previously recorded file system operations such as `open` and `mkdir`, preserving the interarrival time of each system call. We replayed a 30 minute segment of Mummert’s `purcell` trace—this trace captures interactive software development activity such as file editing and compilation. We placed the files accessed by the trace in the Coda distributed file system [10]. During trace replay, the STPM kernel module observed Coda’s network activity as it communicated with a file server to read and write data. We saved the access distribution observed by our module and used it to warm the STPM module before the execution of each scenario.

For each experiment, we measured application performance using the `gettimeofday` system call. To obtain energy measurements, we removed the iPAQ’s batteries and powered the handheld through its external power supply. We sampled the current drawn by the iPAQ approximately 50 times per second using an Agilent 34401A digital multimeter. We calculated system power usage by multiplying each current sample by the mean voltage drawn by the iPAQ—separate voltage samples are not necessary since the variation in volt-



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

age drawn through the external power supply is very small. The total energy consumed by the system over a specific period of time is the sum of the power samples during that period multiplied by the measurement interval.

5.2 Coda distributed file system

We first examined the effectiveness of STPM for Coda file system [10] activity. Coda presents a single file system image across multiple computers by reading file data accessed by a client from remote file servers and writing modifications made on the client back to the servers. Coda caches file data on the client to improve performance and guard against disconnection—on the iPAQ, we used a 16 MB DRAM-based Coda cache. Since Coda uses RPCs for communication, we modified the Coda client to call `TransferHintBegin` before issuing each RPC and `HintEnd` after the RPC completes. The vast majority of Coda RPCs are synchronous—for these RPCs, we issue foreground hints. Although we issue background hints when Coda performs asynchronous writes and prefetches, these activities occur infrequently in our experiments. When Coda starts, it specifies a listen hint that is a fraction of the RPC timeout value—this ensures that server-initiated RPCs arrive at the client promptly. These modifications added 235 lines of source code.

To generate a realistic access pattern, we again used Mumert’s `purcell` trace. From the trace, we selected a different trace segment than the one that we used to warm the STPM access history. The selected trace segment has 10,000 file operations and includes, in total, 42.8 minutes of delay between file operations. These delays typically represent user think time, but can also reflect application processing such as the time for gcc to compile a program. DFSTrace sleeps for the recorded interarrival time before issuing the next file request, so an infinitely fast system would complete trace replay in 42.8 minutes.

5.2.1 Handheld results

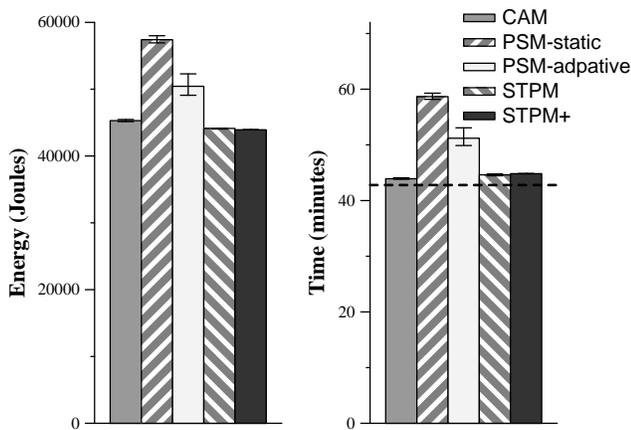
We first replayed the trace using the iPAQ handheld with the Cisco PCMCIA card—Figure 7 shows the results. Since CAM uses no power management, it yields the best possible performance. However, its energy cost is significantly higher than the other strategies. PSM-static (i.e. default 802.11 power management) reduces energy usage but has very poor performance, delaying trace execution by an additional 15 minutes. The dashed line in the rightmost graph shows the total amount of think time in the trace—thus, the cumulative delay added due to file system activity and power management is shown by the portion of each bar that exceeds the dashed line. This delay is especially bad for interactive applications because it represents time that the user must waste waiting for the application making the file accesses to respond. The Cisco card’s PSM-adaptive strategy does substantially better by dynamically switching between PSM and CAM depending upon traffic load. Compared to CAM, PSM-adaptive adds only 6 minutes of cumulative delay and reduces energy usage 14%.

We next replayed the trace on the same hardware using STPM. To provide a fair comparison with card-based power management, we first limited STPM to use only CAM and PSM, i.e. only those modes available to the PSM-adaptive strategy. STPM initially chose to switch to CAM before the 3rd transfer hint in each run and to switch to PSM after 300 ms with no foreground activity. During trace replay, STPM usually became slightly more aggressive and switched to PSM after only 200 ms.

Compared to PSM-static, STPM reduces cumulative delay by 80% and energy usage by 21%. In Figure 7, cumulative delay is compared by examining the height above the dashed line for each bar. Compared to PSM-adaptive, STPM reduces cumulative delay by 58% and energy by 14%. Compared to CAM, STPM adds slightly more than a minute of cumulative delay, but reduces energy consumption by 32%. This means that STPM allows the user to accomplish 48% more work before the battery expires.

How much more energy reduction is feasible? The minimum trace execution time, 44:45 minutes, is given by the time to execute the trace with the card in CAM. The minimum power usage, 1.83 Watts, occurs when the card is in PSM and no data is sent or received. The product of these two values, 4914 Joules, is a lower bound on the minimum energy usage achievable by an adaptive strategy that switches between CAM and PSM. This loose lower bound shows that an optimal strategy could at best be 10% more energy-efficient than STPM. Further, this lower bound is unachievable since it omits the cost of sending data, receiving data, and state transitions—thus, STPM is probably much closer to optimal.

When we allow STPM to disable the Cisco card, it reduces energy usage by an additional 8% as shown by the bars labelled “STPM+” in Figure 7. The mean cumulative delay increases slightly but the difference between STPM and STPM+ is within experimental error. This demonstrates an important advantage of STPM: we did not design an entirely new algorithm to take advantage of the new power management mode. We simply made the benchmark data for disabling the card available to the STPM module and provided a function that performed the relevant state transitions. STPM automatically determined the instances in which the new power mode could be profitably employed.



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

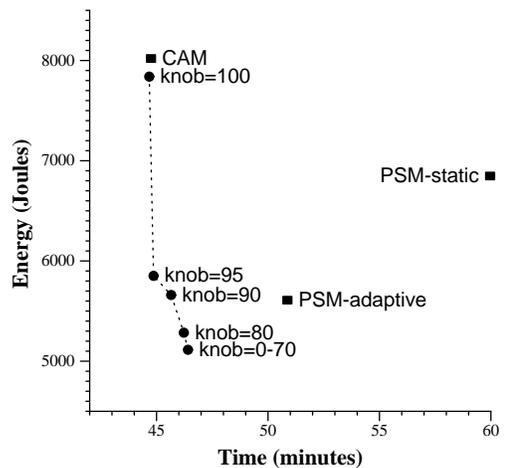
One downside of STPM+ is that disabling the network interface prevents the client from accepting remote connections for ssh, telnet, or similar applications. In contrast, STPM only slightly delays such connections. To support remote connections, STPM+ could periodically poll for incoming connections at an interval specified by a global listen hint. Alternatively, the client could use an external signaling mechanism such as wake-on-wireless [20] to detect new connections and use STPM+ to manage the interface state after connection establishment.

5.2.2 The importance of base power

To demonstrate how well STPM adapts to different hardware platforms, we also replayed the Coda trace on an IBM T20 ThinkPad laptop computer. The base power of the laptop, 15.8 Watts, is an order of magnitude greater than that of the iPAQ.

Figure 8 shows a disturbing result: both PSM-static and PSM-adaptive increase the total energy used to replay the trace! This occurs because the wireless network represents only a small portion of the total power used by the laptop. While switching to PSM reduces the power used by the PCMCIA card by 72%, it only reduces total system power by 6%. Further, PSM-static and PSM-adaptive both increase the time needed to complete the trace because they delay file operations. Although this increase is only about 6 minutes for PSM-adaptive, laptop power usage exceeds 15 Watts during such delays. Thus, as base power increases, the benefit of PSM is less during periods without network activity, and the cost of PSM is higher during periods with frequent network activity. This implies that power management should be more conservative as base power increases.

The energy usage and performance of STPM on the laptop is shown by the fourth bar in each graph in Figure 8. Because STPM accounts for base power, it is considerably more conservative when running on the laptop. On average, it only enters PSM after 15 seconds of inactivity. This strategy works: STPM decreases total system energy usage



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.

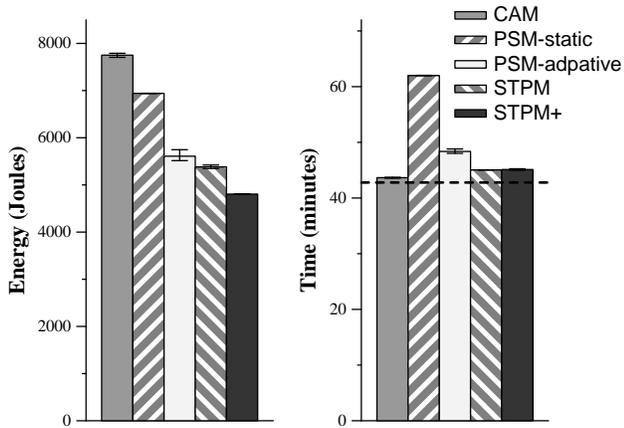
Figure 9: The performance/energy tradeoff

by 2.6% while increasing the total time to execute the trace by only 1.5%. When STPM is allowed to disable the card, results do not change much. The module rarely disables the wireless card because the benefit is so slight—disabling the card reduces system power usage by less than 1% compared to PSM.

5.2.3 Exploiting the performance/energy tradeoff

Using the iPAQ and Cisco card, we explored the impact of changing the performance/energy knob for STPM. Figure 9 plots the tradeoff between energy usage and the execution time for the `purcell` trace. The square marks show the three native modes supported by the Cisco card, while the circles show STPM with five different knob settings. When the knob for STPM is set to 100, the performance is equivalent to CAM (PSM is employed once at the beginning of the trace, accounting for a slightly lower energy usage). The remaining four marks show that as the knob value is decreased, energy conservation improves but performance is decreased. Knob values below 70 yield an equivalent strategy to that realized with a knob value of 0, and hence have the same performance and energy usage.

These results have a powerful property: decreasing the knob value never yields increased energy usage, and increasing the knob value never yields reduced performance. It is clear that the effect of changing the knob value is non-linear. For instance, changing the knob from 100 to 95 substantially reduces energy usage, but changing the knob value from 70 to 0 has no effect. Partly, this occurs because STPM will not choose an inferior strategy if another is available that is expected to yield better performance and energy conservation. For instance, one might expect that STPM would behave identically to PSM-static with a knob value of 0. However, the strategy chosen by STPM yields better performance and greater energy conservation than PSM-static.



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

5.3 Network File System

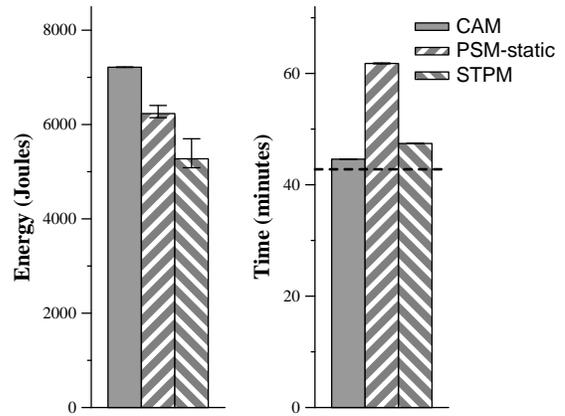
Next, we examined the benefit of STPM for the Network File System (NFS) version 2 [17]. As with Coda, NFS communication is RPC-based—we modified NFS to issue a transfer hint before each RPC begins and end the hint when the RPC completes. Unlike Coda, NFS revalidates cached files before using them—thus, each file access generates at least one RPC. NFS v2 does not delay writing modifications back to the server, nor does it prefetch whole file data. Since there are few sources of background traffic, all hints issued are foreground ones. These modifications added 92 lines of source code.

5.3.1 Cisco results

Figure 10 shows results from replaying the `purcell` trace on the iPAQ using NFS as the underlying distributed file system. For these experiments, we again used the Cisco PCMCIA card. The results for CAM and PSM-static are similar to those for Coda. However, PSM-adaptive does slightly better because NFS issues several concurrent RPCs for large reads and writes. This concurrency generates sufficient network activity for PSM-adaptive to switch immediately to CAM.

During trace execution, STPM becomes more aggressive in switching to CAM. With Coda, STPM switches to CAM before the 3rd RPC; with NFS, STPM switches to CAM before the 1st or 2nd RPC by the end of the trace. During replay, STPM learns that NFS is likely to issue many RPCs closely correlated in time due to its revalidation strategy and a general tendency to issue more RPCs per file operation. STPM gradually adjusts its power management strategy as it observes the NFS request distribution.

Using just CAM and PSM, STPM outperformed PSM-adaptive for NFS, eliminating over 3 minutes of interactive delay and reducing system energy usage by 6%. Using a calculation similar to that in Section 5.2.1, we can derive a loose lower bound of 4313 Joules for energy usage during NFS trace execution—this shows that STPM energy usage is within 11% of optimal.



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

When STPM was allowed to disable the Cisco card, it reduced energy usage by an additional 8% with no noticeable change in performance. Compared to CAM, STPM+ reduced iPAQ energy usage by 38%, which would allow 61% more work on battery power.

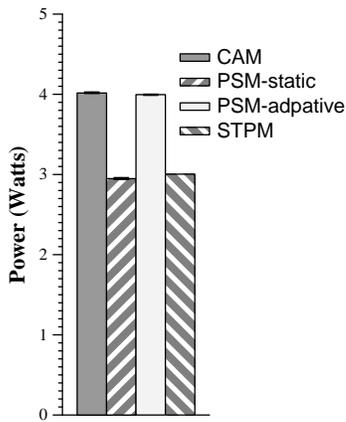
5.3.2 Adapting to interface characteristics

We also wished to verify that STPM adapts to 802.11b cards with significantly different attributes. For this purpose, we selected the Orinoco Silver 802.11b card detailed in Figure 5. Whereas the maximum data rate for the Cisco card is 11 Mb/s, the Orinoco card is limited to 2 Mb/s. Further, the transition costs for the Orinoco card are lower than those for the Cisco card. The Orinoco card cannot be easily disabled to save power, nor does it support an adaptive power management strategy.

Figure 11 shows the results from replaying the `purcell` trace for NFS running on an iPAQ with the Orinoco card. STPM is more aggressive for the Orinoco card—it switches to PSM after only 200 ms compared to a waiting period of 300 ms for the Cisco card. Since the transition costs for the Orinoco card are lower, STPM employs power saving modes more often. STPM significantly outperforms PSM-static, reducing energy usage by 15% and execution time by 23%.

5.4 Xmms streaming audio

We next modified XMMS-embedded to disclose hints when it streams live audio from an Internet server and plays it on an iPAQ. We chose this application because its communication pattern differs significantly from the file systems that we had explored previously. Since the data rate is low (128 Kb/s using the streaming MP3 format), the amount of client buffering is the primary factor that determines whether power management can be used. XMMS-embedded buffers several seconds of audio on the client; thus, it can tolerate the small delays in packet arrival caused by PSM. We modified XMMS-embedded to specify a maximum delay



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

tolerance of 200 ms using `BeginListenHint`. Our modifications added 7 lines of source code.

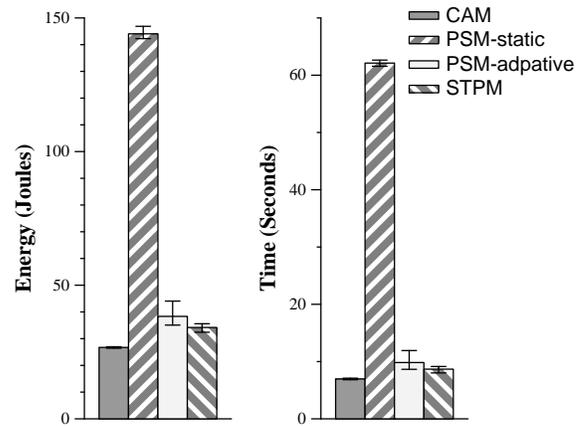
As shown in Figure 12, STPM reduces power usage by 25% compared to CAM. PSM-static is the optimal strategy for this application. However, STPM uses only 2% more energy than PSM-static—this difference reflects the overhead of our module. In contrast, PSM-adaptive always remains in CAM while audio is played. Since PSM-adaptive has no knowledge of application intent, it must assume that latency is critical for audio streaming traffic. In our experiments, no audio packets were dropped because application buffering was sufficient to overcome any jitter caused by power management.

5.5 Thin-client using remote X

Finally, we examined thin-client display of remote X applications on the iPAQ. Remote display allows users to run applications that are too computationally expensive for a handheld, or which have not been ported to the StrongArm. We modified Xfree86 4.3 to give application hints when reading incoming messages and transmitting responses to remote clients. To ensure acceptable interactive response time for long-running remote applications, the X server specifies a maximum delay on incoming traffic of 100 ms whenever a remote session exists. Our modifications added 51 lines of code.

In our experiments, we first started the Gnumeric spreadsheet on the server with its display hosted on a remote iPAQ. We used the iPAQ GUI to load a spreadsheet that contained three columns with roughly 4500 data points each, as well as two complex graphs. We then viewed the spreadsheet and closed the application. We measured the time to perform each action by observing network traffic on the server. We preceded each interactive action with a pause that enabled us to determine when the action began and ended.

Figure 13 shows the cumulative time and energy to perform all interactive activities—it does not reflect any user think time or pause between activities. STPM is over three times more energy efficient and six times faster than PSM-



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

static. STPM uses 12% less energy than PSM-adaptive and is 13% faster. Although CAM mode is 25% faster and uses 28% less energy than STPM, it uses more power during user think time. Figure 14 shows how user think-time impacts energy usage. If think time exceeds 6.5 seconds, STPM uses less energy than CAM over the entire interactive episode.

6. RELATED WORK

To the best of our knowledge, STPM is the first power management algorithm to simultaneously tune its behavior to the characteristics of the network interface, the base power usage of the mobile computer, and the intent and access patterns of applications. STPM differs from previous hint-based approaches to network power management because it requires applications to disclose only current activity, not future activity, and because it uses on-line modeling of application access patterns to set network power management parameters.

Lu et al. have also explored the use of application hints in power management [14]. They use a predictive policy to estimate when to wake up the network interface. Their hints are at a coarser granularity than those used by STPM and require applications to have knowledge of future activity: for instance, they allow an editor to specify that it requires the network interface over the next few minutes to auto-save a buffer. Their work focuses on scheduling network availability for applications with relatively flexible deadlines, while STPM focuses on supporting latency-sensitive applications. From a broader perspective, application hints have also been successfully used to improve disk power management [24] and I/O prefetching [15, 19].

Kravets and Krishnan [12] also advocate an application level power management strategy. Their work provides a transport level protocol to improve energy usage. STPM complements their solution in that it predicts opportunities for performance improvement and energy conservation.

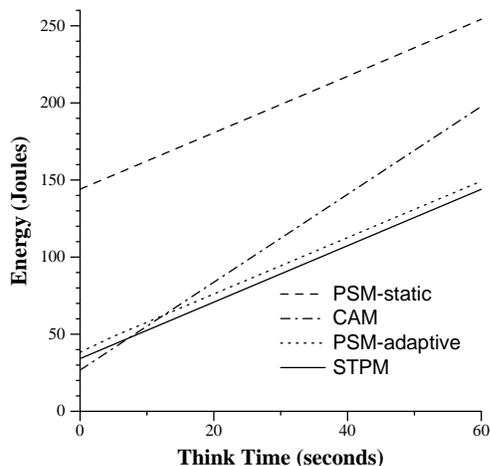


Figure 14: Effect of think time on X energy usage

Wake-on-wireless [20] uses a low-power network to signal a mobile client when packets are waiting at the base station. Current wireless network transition costs are prohibitive for using this technique to disable and enable the network card between transfers for applications like NFS. However, wake-on-wireless seems extremely promising for supporting incoming connections for server processes like the X server and sshd. We can envision a hybrid strategy that uses wake-on-wireless to detect incoming connections during idle periods and uses STPM to manage existing connections.

Krashinsky’s bounded slowdown (BSD) protocol [11] disables the network interface to save power while bounding the relative delay on transfer round-trip time. BSD differs from STPM in that it does not explicitly consider transition costs or the base power usage of the device. Unlike STPM, BSD is implemented without knowledge of application intent. The advantage of not using application hints is that BSD can show benefit for unmodified applications. The disadvantage is that BSD must behave conservatively, limiting the potential for energy savings. Further, BSD requires some small modifications to the 802.11b protocol, making deployment on current hardware more difficult.

Simunic’s TISMDP [21] operates at a coarser granularity than STPM. TISMDP only decides when to transition from a higher-power state to a lower-power state. It immediately resumes upon the arrival of the next request. Thus, for intermittent activity like NFS RPCs, a transition cost must be paid for each RPC.

Chandra [2] explores how the regular nature of streaming multimedia can be exploited to improve power management. Stemm et al. [22] investigate methods that reduce energy consumption of network interfaces for electronic mail and web browsing applications on PDAs. Unlike STPM, both these techniques benefit only a limited set of specific applications.

Adaptive methods have also been applied to disk [4, 7] and CPU [5] power management. Though these domains differ substantially from wireless network power management, we believe that many of the design principles that we have used for STPM can be applied.

7. CONCLUSION

Wireless network power management can severely degrade the performance of latency-sensitive applications and even increase the total energy needed to perform interactive activities on a mobile computer. In order to provide significant energy conservation without substantial performance degradation, a power management strategy should be tuned to reflect application intent and access patterns, as well as the power characteristics of the network interface and mobile computer. It is infeasible to expect a user to manually tune the power management algorithm for each combination of application, network interface, and mobile computer. Therefore, we have built a *self-tuning* power management module that adapts its behavior in response to a changing environment. Our results show that self-tuning improves both performance and energy conservation compared to current power management strategies.

Our Linux implementation of self-tuning power management offers an important opportunity for future work. We plan to deploy our module to a small user community and gather detailed feedback about the benefits of self-tuning power management. This deployment will require us to both broaden the set of applications that disclose power management hints and provide support for non-hinting applications. We also plan to adapt STPM to the management of other types of wireless network interfaces and different system components such as hard disk drives.

Acknowledgments

We would like to thank Mark Corner, Brian Noble, and the anonymous reviewers for several helpful suggestions that improved the quality of this paper. This research was partially supported by the National Science Foundation under grant CCR-0306251 and by an equipment grant from Intel Corporation. The views and conclusions contained in this document are those of the authors and should not be interpreted as representing the official policies, either expressed or implied, of NSF, Intel, the University of Michigan, or the U.S. government.

8. REFERENCES

- [1] BOX, D., EHNEBUSKE, D., KAKIVAYA, G., LAYMAN, A., MENDELSON, N., NIELSEN, H. F., THATTE, S., AND WINER, D. Simple Object Access Protocol (SOAP) 1.1. Tech. rep., W3C Note, May 2000.
- [2] CHANDRA, S. Wireless network interface energy consumption implications of popular streaming formats. In *Proceedings of Multimedia Computing and Networking* (San Jose, CA, 2002), pp. 85–99.
- [3] CISCO SYSTEMS, INC. *Cisco Aironet wireless LAN client adapters installation and configuration guide for Linux*.
- [4] DOUGLIS, F., KRISHNAN, P., AND BERSHAD, B. Adaptive disk spin-down policies for mobile computers. In *Proceedings of the 2nd USENIX Symposium on Mobile and Location-Independent Computing* (Ann Arbor, MI, April 1995), pp. 121–137.
- [5] FLAUTNER, K., REINHARDT, S., AND MUDGE, T. Dynamic voltage scaling on a low-power microprocessor. In *Proceedings of the 7th Annual International Conference on Mobile Computing and*

- Networking (MOBICOM '01)* (Rome, Italy, July 2001), pp. 260–271.
- [6] FLINN, J., AND SATYANARAYANAN, M. Energy-aware adaptation for mobile applications. In *Proceedings of the 17th ACM Symposium on Operating Systems Principles (SOSP)* (Kiawah Island, SC, December 1999), pp. 48–63.
- [7] GREENAWALT, P. Modeling power management for hard disks. In *Proceedings of the Symposium on Modeling and Simulation of Computer Telecommunication Systems* (Durham, NC, January 1994), pp. 62–66.
- [8] HOWARD, J. H., KAZAR, M. L., MENEES, S. G., NICHOLS, D. A., SATYANARAYANAN, M., SIDEBOTHAM, R. N., AND WEST, M. J. Scale and performance in a distributed file system. *ACM Transactions on Computer Systems* 6, 1 (February 1998).
- [9] IEEE LOCAL AND METROPOLITAN AREA NETWORK STANDARDS COMMITTEE. *Wireless LAN medium access control (MAC) and physical layer (PHY) specifications*. New York, New York, 1997. IEEE Std 802.11-1997.
- [10] KISTLER, J. J., AND SATYANARAYANAN, M. Disconnected operation in the Coda file system. *ACM Transactions on Computer Systems* 10, 1 (February 1992).
- [11] KRASHINSKY, R., AND BALAKRISHNAN, H. Minimizing energy for wireless web access with bounded slowdown. In *Proceedings of the 8th Annual International Conference on Mobile Computing and Networking (MOBICOM '02)* (Atlanta, GA, July 2002).
- [12] KRAVETS, R., AND KRISHNAN, P. Application-driven power management for mobile communication. *ACM Wireless Networks* 6, 4 (2000), 263–277.
- [13] LORCH, J. R., AND SMITH, A. J. Improving dynamic voltage scaling algorithms with PACE. In *Proceedings of ACM SIGMETRICS* (Cambridge, MA, June 2001).
- [14] LU, Y.-H., BENINI, L., AND MICHELI, G. D. Power-aware operating systems for interactive systems. *IEEE Transactions on VLSI* 10, 2 (April 2002), 119–134.
- [15] MOWRY, T. C., DEMKE, A. K., AND KRIEGER, O. Automatic compiler-inserted I/O prefetching for out-of-core applications. In *Proceedings of the 2nd USENIX Symposium on Operating Systems Design and Implementation* (Seattle, WA, October 1996), pp. 3–17.
- [16] MUMMERT, L., EBLING, M., AND SATYANARAYANAN, M. Exploiting weak connectivity in mobile file access. In *Proceedings of the 15th ACM Symp. on Op. Syst. Principles* (Copper Mountain, CO, Dec. 1995).
- [17] NETWORK WORKING GROUP. *NFS: Network File System protocol specification*, March 1989. RFC 1094.
- [18] NYE, A., Ed. *X Protocol Reference Manual*. O'Reilly and Associates, Inc., 1990.
- [19] PATTERSON, R., GIBSON, G., GINTING, E., STODOLSKY, D., AND ZELENKA, J. Informed prefetching and caching. In *Proceedings of the 15th ACM Symp. on Op. Syst. Principles* (Copper Mountain, CO, Dec. 1995).
- [20] SHIH, E., BAHL, P., AND SINCLAIR, M. J. Wake on wireless: An event-driven energy saving strategy for battery operated devices. In *Proceedings of the 8th Annual International Conference on Mobile Computing and Networking (MOBICOM '02)* (Atlanta, GA, July 2002).
- [21] SIMUNIC, T., BENINI, L., AND DE MICHELI, G. Energy-efficient design of battery-powered embedded systems. In *Proceedings of the 1999 International Symposium on Low Power Electronics and Design* (November 1999), pp. 212–217.
- [22] STEMM, M., AND KATZ, R. H. Measuring and reducing energy consumption of network interfaces in hand-held devices. *IEICE Transactions on Fundamentals of Electronics, Communications, and Computer Science* 80, 8 (August 1997), 1125–1131.
- [23] SUN MICROSYSTEMS, INC. *Java Remote Method Invocation specification*, December 1999.
- [24] WEISSEL, A., BEUTEL, B., AND BELLOSA, F. Cooperative I/O: A novel I/O semantics for energy-aware applications. In *Proceedings of the 5th Symposium on Operating Systems Design and Implementation* (Boston, MA, December 2002), pp. 117–129.
- [25] ZENG, H., ELLIS, C. S., LEBECK, A. R., AND VAHDAT, A. ECOSystem: Managing energy as a first class operating system resource. In *Proceedings of the 10th International Conference on Architectural Support for Programming Languages and Operating Systems (ASPLOS-X)* (San Jose, CA, October 2002).