Scheduled PSM for Minimizing Energy in Wireless LANs

Yong He[1], Ruixi Yuan[1], Xiaojun Ma[2], Jun Li[2], Charles Wang[2]

[1] Dept. of Automation, Tsinghua University, Beijing, 100084, China. heyong04@mails.tsinghua.edu.cn, ryuan@tsinghua.edu.cn.

[2] Thomson Broadband, R&D, Co., Ltd, Beijing, 100085, China. {Xiao-jun.ma, Jun.li, Charles.wang}@thomson.net.

Abstract*—Power conservation is a general concern for mobile computing and communication. In this paper, we investigate the performance of the current 802.11 power saving mechanism (unscheduled PSM) and demonstrate that background network traffic can have a significant impact on the power consumption of mobile stations. To improve power efficiency, a scheduled PSM protocol based on time slicing is proposed in this paper. The protocol adopts the mechanism of time division, schedules the access point to deliver pending data at designated time slices, and adaptively adjusts the power state of the mobile stations. The proposed scheme is near theoretical optimal for power saving in the sense that it greatly reduces the effect of background traffic, minimizes the station idle time, and maximizes its energy utilization. Comprehensive analysis and simulations are conducted to evaluate the new protocol. Our results show that it provides significant energy saving over the unscheduled PSM, particularly in circumstances where multiple traffic streams coexist in a network. Moreover, it achieves the saving at the cost of only a slight degradation of the one-way delay performance.

I. Introduction

The proliferation of portable computing and mobile technology has posed a great concern on energy conservation. A major constraint for many wireless applications is the limited size and lifetime of the batteries that power up these mobile computing devices. Without careful design of an energy-aware system, or typically a power management module on the wireless interface, a mobile device's energy can be drained out quickly by the upper layer applications. Our experiments on a HP iPAQ hx2750 PDA showed that an 802.11b wireless adapter with power management disabled can shorten its lifetime up to 50% for light to moderate traffic loads. In this paper, we seek to minimize the energy consumed by a wireless network interface running IEEE 802.11 protocol, and particularly focuses on the wireless LANs and infrastructure based network architecture, a currently dominating network paradigm in home, office environments and public hotspots.

As specified in the standard, an IEEE 802.11 based wireless network interface can choose to stay in one of two states at any moment, awake or sleep. In the awake state, the radio is powered up and the wireless interface can perform data transmission or reception, or stay in idle and wait. In the sleep state, on the contrast, the radio is turned off and the wireless interface cannot detect or sense the network behaviors of others. Wireless interface in awake state usually consumes an order of magnitude more power than that in sleep state. The major task of power management is to choose proper time and sequence of the state transition between the two states. A good solution should wake up the wireless interface at the proper time for communication, and let it enter into sleep state to save energy if no traffic is direct to it.

A power saving mode (PSM) is defined in the 802.11 specification [1] for power management. In this mode, the access point (AP) buffers incoming frames destined for mobile stations in PSM and periodically announces its buffering status through the traffic indication map (TIM) contained in the beacon frames. The mobile station wakes up periodically to listen to the beacon frames. In the unicast case, once the bit corresponding to its association ID (AID) is set in the TIM, the mobile station initializes a PS-Poll frame to the AP to retrieve data and the AP responds each poll with one buffered frame. Multiple PS-Polls are allowed until all outstanding frames have been retrieved. In the broadcast/multicast (B/M) case, the existence of buffered B/M frames is indicated by setting the B/M traffic indication bit in the delivery TIM (DTIM), which is a special TIM sent out at a fixed number of beacon intervals. All B/M frames buffered at the AP are delivered immediately after the beacon frame containing DTIM. As opposed to the normal continuous active mode (CAM), a mobile station in PSM can often have opportunities to turn its network interface off to save energy, given that it has no data pending at the AP. For light to moderate traffic load, the legacy PSM can greatly reduce the energy consumption and extends the mobile stations' lifetime.

However, the legacy PSM becomes inefficient when multiple traffic streams coexist in a network. Our analysis and experiment results in section II show that the contention of medium access among multiple stations increases the deferring time for delivery of buffered frames, and thus the time to enter to sleep mode is delayed. The power consumption of a mobile station depends not only on the traffic load destined to it, but also on the background traffic from other mobile stations. A heavy background traffic load often causes high power consumption. This phenomenon is due to three factors. First, the one-poll-one-response manner to retrieve data is not efficient, and further it is susceptible to be interrupted by others' network activities, which leads to longer wakeup time for the mobile stations. Second, though a station can be informed of the existence of buffered data, it has no way to

1-4244-1588-8/07/$25.00

know the starting and ending point of the delivery procedure. It has to stand by and passively wait for the conclusion of data delivery. Third, the legacy PSM does not differentiate multicast and broadcast traffic. A mobile station has to stay in active state to receive B/M data if the B/M bit in the DTIM is set, even it does not join any multicast group.

This paper aims to design a solution that improves the power performance of the legacy PSM, diminishes the impact of background traffic on power consumption, while not degrading its networking performance too much. Analysis in section II indicates that a protocol with scheduled, uninterrupted delivery procedure and differentiated multicasting is a preferable choice to achieve our goal. In section III, a time slicing base PSM is proposed. In this scheme, a beacon period is divided into multiple time slices, and all pending frames are scheduled to deliver at the appointed time slices. The TIM structure is redesigned to convey the slice assignment information. A mobile station can learn from the TIM not only whether or not it has data pending at the AP, but also the timing of the expected data delivery procedure. Hence it is possible for the mobile station to control the state of the wireless interface adaptively to the traffic arrangements. During the time slices that are assigned to other mobile stations, it can choose to enter into sleep state to save energy. As the proposed scheme use scheduled transmission for delivery of pending data, we name it as scheduled PSM in this paper. For comparison, the legacy PSM is called unscheduled PSM or PSM for short.

In section IV, comprehensive simulations are conducted on *ns2* to study the performance of scheduled PSM, compared with the unscheduled PSM. The simulation results show that the scheduled PSM provides significant energy saving over the unscheduled PSM when multiple traffic streams coexist in the network. In all cases with fixed data rate of foreground traffic, its power consumption remains stable when data rate of the background traffic load increases, while that of the unscheduled PSM exhibits a distinct upward trend. Moreover, it achieves this at cost of only a slight degradation of the one-way-delay performance. Our simulation results show that, the one-way-delay of both uplink and downlink streams is increased by less than 15ms compared to the unscheduled PSM, within an acceptable level for most applications.

In section V, we describe several newly proposed power saving mechanisms in 802.11 standard groups (802.11e, 802.11n and 802.11v), discuss their relative merits and disadvantages, as well as presenting some related work in the literature. Section VI gives a brief summary.

II. IMPACT OF BACKGROUND TRAFFICS ON UNSCHEDULED PSM

In this section, we investigate the impact of the background traffic on the power performance of the unscheduled PSM. We have set up a small wireless LAN with commercially available products: a Dell D600 laptop serving as a server, two HP iPAQ hx2750 PDAs with IEEE 802.11b build-in Wi-Fi adapters serving as clients, and a DWL-2100AP serving as the AP. One of the PDAs is chosen as the measuring point, and its traffic

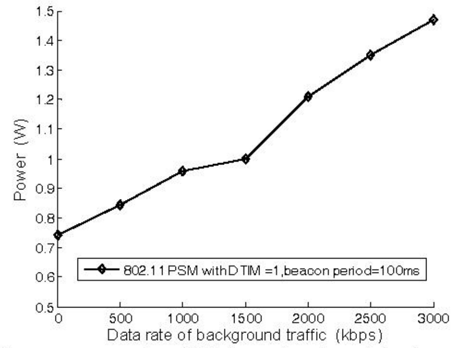

Fig. 1 Power consumption of PDA as a function of the data rate of the background multicast traffic. No foreground traffic is imposed.

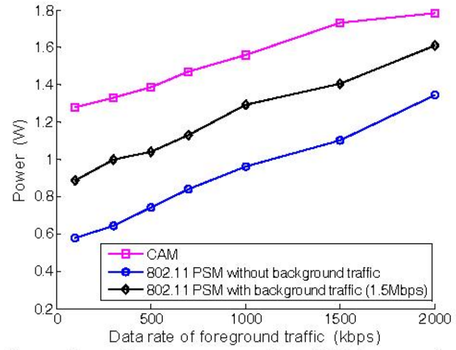

Fig. 2 Comparison of power consumption of PDA among three cases: CAM, PSM, and PSM with background traffic.

with the AP is viewed as the foreground traffic, while traffic destined to the other client is viewed as the background traffic. An Ixchariot [15] tool is used to generate traffics between clients and server. The power consumption of the PDA is measured through a system API *GetSystemPowerStatusEx2()*. Its precision reaches 1mW and its value updates every 5 seconds. Each measurement in our experiment lasts over 30 minutes and each data point is an averaged of 10 measurements. Two representative results are shown in Fig. 1 and Fig. 2.

Fig. 1 shows the impact of the background multicast traffic on the power consumption of the PDA. Though no foreground traffic is imposed on the client PDA, its power consumption exhibits a sharp upward trend with the increasing data rate of the background traffic. In Fig. 2, the power consumption as a function of the foreground traffic is compared among three cases: CAM, PSM without background traffic and PSM with a 1.5Mbps background multicast traffic. The gap between these curves demonstrates that the unscheduled PSM does save power, but the saving is reduced by the presence of background traffic.

The reduction in power saving is caused by the following factors. First, the AP does inform the existence of the buffered data for a station in the TIM structure, but it does not give any information of the starting and ending time of the delivery of the data. Thus before all pending data are retrieved, the mobile has to stay awake to listen to all frames in the air so as not to miss any data destined for it. Second, the one-poll-one-response mechanism to deliver unicast frame is not only inefficient for a mobile station to retrieve its buffered frames at the AP, but also is susceptible to be interrupted by transmissions of other stations and leads to a prolonged retrieving period. Third, using only one bit to indicate the traffic state of broadcast/multicast traffic at AP makes frames belonging to broadcast and different multicast groups undistinguishable. As illustrated in Fig. 1, a mobile station with no traffic has to pay extra energy for receiving multicast frames not destined to itself.

An efficient power saving mechanism should address the three factors above. In section III, we design a new PSM that employs scheduled, uninterrupted frame delivery procedure and also differentiates broadcast and multicast traffics.

III. SCHEDULED PSM

To mitigate the impact of the background traffics on the power performance of the unscheduled PSM, a scheduled PSM is proposed and described in this section. It is based on time slicing and uses scheduled transmission for delivery of buffered frames. It is motivated, in part, by the time division multiplexing technique employed by DVB-H and derived, in part, from our previous experiments and analysis in section II. The main idea is to divide a beacon period into multiple slices and schedule the pending frames for delivery.

A. Overview of the proposed scheme

In the scheduled PSM, the AP works as the central coordinator and scheduler during the entire procedure. It buffers all incoming data frames destined to the mobile stations. Each time before transmitting a beacon, the AP divides the beacon period into a fixed number of time slices of equal length. The number has the form of 2^n-1, with n being the number of bits used in the TIM element to identify a time slice. After the division, the AP assigns slices for the buffered data frames. Contiguous time slices are assigned to the same destination address for transmission. After that, the AP assembles traffic indicators and slice arrangements into a TIM element, and broadcast it to mobiles stations in the beacon. In our scheme, The TIM structure is redesigned to accommodate the slice assignment information for the scheduled PSM.

Note that a time slice is assigned to a specific destination address instead of a mobile station, as in the scheduled PSM, the AP treats the pending broadcast/multicast (B/M) frames as the same with the unicast frames. A multicast group is assigned with a unique association ID (AID) just as do for a mobile station. And AID 0 is dedicated to the broadcast. Through the multicast association, the AP establishes a mapping between destination addresses and the AIDs, and each frame in the buffer is identified by its corresponding AID. Frames with identical AID are grouped together and handled in batch, that is, they are delivered in consecutive time slices, forming a service period (SP) for that group or AID.

The AP employs a scheduling algorithm to calculate the number of time slices needed for each group. Uplink traffic and frame error rate should be considered in the algorithm. The AP is responsible for arranging the time slices in an appropriate sequence and expressing it semantically through the TIM element. The AP should be aware of the slice arrangement all the time. The AP initiates frame transmissions to deliver the pending data a destination when the scheduled SP starts. The AP must end the frame transmission for the particular destination before the beginning of next SP, thus avoiding interference with the transmission for the next SP. In other words, the SP is actually a time limit imposed on the frame transmission period for a specific destination. As an example, Fig. 3 shows the time slice assignment in a beacon period. The beacon period is divided into 15 equal-length time slices, and 13 out of them forms 4 consecutive SPs and assigned to destinations with AID 0, 4, 1 and 2 separately. The other 2 slices are unassigned and reserved for normal DCF. Each slice lasts for $aSliceDuration = tBeaconPeriod/nSlice = 100/15 = 6.667ms$ for a typical 100ms beacon period. The AP uses the value of *aSliceDuration* and slicing index (SI), which identifies the first time slice assigned to a SP, to calculate the starting and ending time of that SP. For example, consider the third SP corresponding to AID 1, the AP should initiate transmission of data for AID 1 at the time of $(SI(1)-1) \cdot aSliceDuration = (7-1) \times 6.667 = 40\ ms$, and conclude before the time of $(SI(2)-1) \cdot aSliceDuration = (11-1) \times 6.667 = 66.67\ ms$.

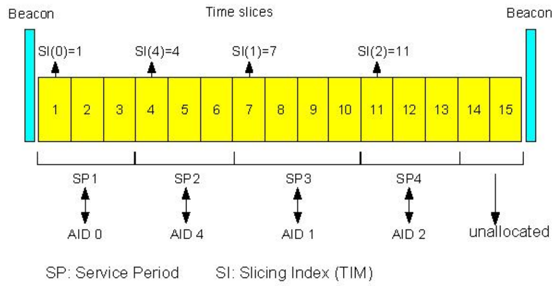

Fig. 3 A beacon period with 15 time slices.

Each mobile station maintains an AID list. It consists of one unicast AID assigned during the association procedure, and one or more multicast AIDs assigned during multicast association procedure if the mobile station joined multicast groups. Still there is an AID 0 in the list, used for broadcast and known by everyone. The AID list serves as a filter to pick out the interested bits from the partial virtual bitmap (PVB) in the TIM element. Every time a beacon frame is received, the mobile station checks the bitmap using the AID list, and if a bit corresponding to an AID in the list is set, the station knows that there are data pending at the AP for this AID. To find out the assigned time slices, it has to go further towards the slicing map field and looks up the slicing index (SI) that matches the AID. The value of SI indicates the first time slice assigned for that AID in the slice arrangement. It is used as a parameter to derive the starting time of the service period for the AID by the mobile station. The length of the service period can be obtained by finding the next higher SI in the slicing map. Such a bit-checking procedure repeats until all bits corresponding to the AID list have been checked out.

The mobile station switches its power state of network interface between sleep and awake to the arrangement of time slices in the TIM element. It enters into sleep state if it does not interest in the data of the current SP, and wakes up until the next interested SP. For example, it chooses to sleep during all multicast SPs if it does not join any multicast group. If the saved energy during sleep does not counteract against the extra energy used for the state transition, it chooses to remain in awake state. Whenever the mobile station wants to send data, it wakes up and sends them out using the normal DCF. If the data are not time sensitive, their transmissions are postponed to the next wakeup period.

A requirement for the scheduled PSM is that the AP and the mobile station should agree on the starting time of a SP, otherwise the mobile station may miss the desired SP. As the time is calculated relatively to the transmission of beacon frames, the disagreement can be relieved if proper events are selected as the benchmarks for time calculation. In the proposed scheme, the signal generated at the PHY layer to confirm the reception (for mobile stations) or transmission (for AP) of a beacon frame is used as the benchmark. The time discrepancy of the occurrence of such events between AP and mobile stations is often several microseconds. In our subsequent analysis and simulation, we assume that the mobile station wakes up *500us* earlier than the expected time. This early wakeup does cost some energy, but this overhead is unavoidable and is minimal compare to the overall energy consumption.

B. Redesign of the TIM structure

The proposed solution requires that more information should be piggybacked onto the beacons to inform mobile stations when to wake up and when to sleep. The original 802.11 TIM structure can not fulfill this task, thus a new TIM structure is designed in scheduled PSM, as depicted in Fig. 4.

Compared with the original TIM structure, two fields, slicing control field and slicing map field, are appended to the end of the partial virtual bitmap (PVB) in order to accommodate the slice arrangement information. In addition, the significance of each bit of the PVB is extended to support multicast. The single B/M bit is still in use but restricted to broadcast. The detailed changes are described as follows.

Slicing control field (newly added)

This field regulates the number of bits used to represent the slicing index (SI) in the slicing map field. Meanwhile, it is used to derive the number of time slices that should be divided up within one beacon period. For a slicing control field with a value of n, the beacon period is divided into 2^n-1 time slices, each lasting for a period of $tBeaconPediod/(2^n-1)$.

The value of the slicing control field can be changed dynamically in different network environments and in different traffic conditions. Higher value of this field means finer granularity of the time division and more elaborate arrangement of slice assignment thus yields more energy conservation, but the cost is longer slicing map field within the beacon. Hence setting an appropriate value in this field is a trade-off between power performance and the burden of beacon frames. Our experience is that a value of 5 to 8 can handle most

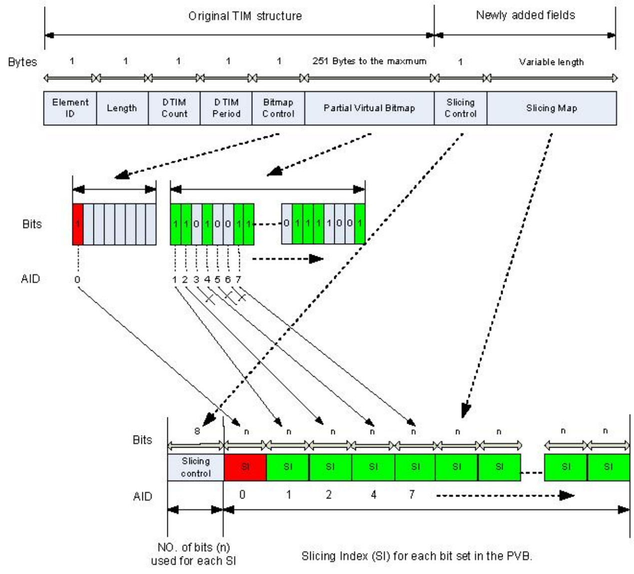

Fig. 4 Structure of the new TIM element.

situations. Moreover, this value should be set to ensure that at least a nominal-sized frame can be transmitted in a time slice.

Slicing map field (newly added)

The slicing map field consists of a string of slicing indexes (SIs). The value of SI is an integer represents the sequence number of the first time slice of a SP in the slice queue. It can be used to calculate the starting time of the SP in a beacon period for a specific destination. To save bits in slicing map field, only a bit that is set to 1 in the bitmap (including the first bit in Bitmap control field) would have a SI in the slicing map field corresponding to it. These SIs are arranged in an order in compliance with the relative position of the corresponding bits in the bitmap, or in other words, they are arranged so that their AIDs are in an ascending order. The omission of SIs in the slicing map whose corresponding bits are cleared in the bitmap significantly reduces the overhead of the beacon frame.

Though the value of a SI is usually a positive integer, it is possible that a SI is filled a value of zero, indicating that there is no more slices available in the current beacon period for the AID corresponding to this SI. A mobile station received such a message would wait for the AP to deliver its pending data in the next beacon interval, or voluntarily issues PS-Polls to the AP to retrieve its data as does in the unscheduled PSM.

Partial virtual bitmap (PVB) field (changed)

The semantics of bits in the in PVB has been extended to represent the traffic status of the multicast groups. Through the multicast association procedure, a multicast group is assigned a unique AID, and the AP can use this AID to allocate a bit from the bitmap as a traffic indicator dedicated to that multicast group, as does in the unicast case.

The B/M traffic indicator (changed)

The significance of the one-bit B/M traffic indicator in the original TIM is now restricted to represent only the status of broadcasting traffic in the new TIM structure.

The new TIM structure is the key of the design of the scheduled PSM. It carries all needed information for the operations of the scheduled PSM. Using the new TIM, the AP can negotiate with mobile stations to fulfill the scheduled delivery procedure. Yet, the increased overhead on the beacon frame is small, given that the value of slicing control field is properly chosen. For example, if the value of slicing control field is 6 (sufficient for most situations) and there are *m* bits set in the bitmap, then the additional cost of overhead is *(1 + ceil((m * 6)/8))* bytes (*ceil(x)* is a function that takes the next higher integer than x). For m = 8, the cost is merely 7 bytes.

Another important aspect of the new TIM design is the consideration of backward compatibility. If the semantics of the B/M indicator bit remains the same as that in the unscheduled PSM, then the new TIM can support the legacy mobile stations and the schedule PSM can be applied to the legacy 802.11 networks. The cost is that the AP has to reserve a SI dedicated to broadcast in the slicing map field and a value zero of the SI indicates the non-existence of broadcast data, while a positive value indicates the starting time slice for broadcast data. For a mobile station running scheduled PSM, it can see the sliced frame transmission, whereas a mobile station running legacy PSM is not aware of this. By this modification, the scheduled PSM is completely backward compatible.

C. Multicast association/disassociation procedure

The multicast association procedure works in a similar manner to the (re)association procedure defined in the original 802.11 standard. Once a station has been confirmed by the router (through IGMP messages) that it has joined a multicast group, it initiates a multicast association request to the AP, with some field containing the MAC address of the multicast group. The AP responds it with a multicast association response frame filled with an AID for that MAC address. The AP is responsible for assigning AIDs to the new multicast addresses, and maintaining consistency of AIDs for the same multicast address among mobile stations.

When a station decides to quit a multicast group, it sends a multicast disassociation request to the AP with one field containing the multicast MAC address of the group. The AP responds with a disassociation response and withdraws the resources not used any longer or updates the relevant state. For example, if the disassociated mobile station is the last member in the multicast group, then after the successful disassociation, the allocated AID and the corresponding bit in the bitmap can be taken back and reused by others.

D. Service period (SP)

A SP is a contiguous time during which the buffered data at the AP are transmitted. It consists of one or more consecutive time slices, with the starting and ending slices are indicated by the corresponding SIs in the slicing map. It is designed for a specific destination to retrieve pending data from AP. However, it does not mean that it is completely dedicated to that destination. In fact, to accommodate the uplink traffic and frame retransmissions, the AP should allocate more time slices than needed for a destination. Moreover, it may choose to cease its transmission of data in advance to reserve a portion of time in the SP for the normal DCF. Thus an SP is consisted of scheduled period used for delivery of pending data and contention period reserved for normal frame transmissions. The proportion between the two is decided by the AP.

In the scheduled PSM, each SP is initiated by the AP. To save energy of mobile stations, the AP should try to start a SP at the schedule time as accurate as possible. To avoid contention with mobile stations, the AP gains control of the medium after the wireless medium has been sensed to be idle for a PIFS time. As the PIFS is shorter than the normal DIFS, the AP takes a higher priority over the mobile stations to access the medium. However, even by this approach, it is still possible that the SP is delayed by the transmission of uplink traffic.

The AP uses protection mechanism for its frame transmission in a SP. If the desired destination is an individual mobile station, the AP uses the RTS/CTS frame exchange to start a new SP; otherwise, the Self-CTS protection is used. To provide uninterrupted delivery of pending data, the transmission procedure operates in a bursty manner, just as the transmission procedure for fragmented frames in DCF and for the TXOPs in 802.11e [2]. That is, a new transmission of frame commences at SIFS later after the completion of the immediately preceding frame exchange sequence. Moreover, to avoid wasting the medium resource, the AP updates the network's NAV settings frame by frame, which means that the duration in a transmitting frame only covers the next frame exchange sequence (if any).

The total time to deliver all the pending data for a destination may exceed the length of the allocated SP for it, particularly when retransmissions occur frequently. In this case, the delivery procedure must cease before expiration of the current SP and the AP indicates the existence of more buffered data by setting the *MoreData* bit in the frame header to 1. A mobile station receives this message can initiatively poll the remaining data from AP subsequently if the data is unicast; or wait for the AP to deliver its data in the next beacon period.

IV. Performance Evaluation

Comprehensive simulations have been conducted to study the power performance of the scheduled PSM, compared with the unscheduled PSM and the theoretical optimal performance. We emphasized on the impact of background network traffics on the power performance of individual mobile stations. Furthermore, we also evaluated the network performance of the two power saving schemes in terms of one-way traffic delay.

A. Simulation methodology

A new power saving module is implemented on the network simulator *ns2* [16] to simulate both the unscheduled and scheduled PSM. In this module, two buffering queues are maintained at the AP. One is a sending queue that holds all frames destined to mobile stations in active stations, and the other is a power saving queue used to accommodate the frames destined for power saving stations. A frame is dequeued from the power saving queue to the sending queue if it is ready to send. The AP maintains status and information for each associated mobile stations that are necessary for power management. The timing of state transition of network interface is controlled by a scheduler on the MAC layer, which periodically queries the power state of the PHY layer and collects information to make decision. In addition, the multicast association/disassociation procedure and beaconing with new TIM structure is implemented as management entities in the MAC layer, and an address list that filters the incoming frames is place on the MAC layer to support multicast.

To simulate the scheduled PSM, a simple algorithm for assigning slices for mobile stations is designed. Before transmission of each beacon, the total time for delivery of the buffered frames for each destination is calculated based on the size of these frames, the nominal data rate and the frame error rate. The time includes the overheads of the control frames. For example, for an n-byte unicast frame, the expected time to complete its transmission is (*aPLCPPreambleTime+n·8/DataRate+ aSIFSTime+ aPLCPPreambleTime + ACKsize/ACKrate) · f(FrameErrorRate),* where *f(FrameErrorRate)* is a function of the frame error rate and it accounts for the increased time due to retransmission. Once the expected transmission time (T) has been derived, the scheduler of the AP can determine the number of slices (N) that should be allocated to the destination, using the additional parameters: beacon period (D), number of slices within one beacon period (M), and surplus bandwidth allowance (A). The value of A specifies the excess allocation of time for the normal DCF within the allocated service period.

$$N = \lceil A \bullet M \bullet T / D \rceil \qquad (1)$$

In the current implementation, we let *f(FrameErrorRate)= 1/FrameErrorRate*, and set A to 1.40 with the frame error rate 10%, and M is set to 64 for a beacon period of 100ms.

To simulate the power consumption of the 802.11b interface card, we used an energy model that derived from previous work [5]. The consumed power is set to 1.346W while transmitting, 0.9W while receiving, 0.741W while idle, and 0.048W while sleeping. To better simulate the activities of the wireless interface card, we modeled the power usage of the transition between idle and doze state as 0.6mJ, and the transition time is set to 0.4ms. The introduction of energy consumption in state switching may affect the station's decision of the state transition of wireless interfaces.

The simulation modeled a network consisting of an AP and three mobile stations. The AP is a special station whose upper layer can work as both traffic generators and traffic sinks. It can server as a data server and streams can be created between the AP and mobile stations directly. Stations including AP are numbered from 0 to 3. Station 0 is the AP, and station 1 is the measuring point throughout our simulations. The traffic imposed between AP and station 1 serves as the foreground traffic, while traffic imposed between AP and station 2 and 3 are served as background traffic. The radio range is set large enough to guarantee the connectivity between any two stations. All traffics are constant bit rate (CBR) UDP flows, with the packet size fixed to 1000 Bytes. Unless explicitly point out, the traffics are downlink.

TABLE I
SIMULATION PARAMETERS

Parameters	Value
Beacon Period	100ms
DTIM Period	3
Listen interval	1
Minimal contention window	31
Maximal contention window	1023
Slot time	20us
SIFS	10us
Preamble duration	144us

We chose 802.11b as the default wireless protocol, and set both the data rate and the basic rate to 11Mbps. The RTS threshold is set to 0 to enable the RTS/CTS protection. The setting of other parameters of 802.11b is shown in table I. Each simulation run lasts for 200 seconds, and each data point is an average of 10 simulation runs.

B. Power performance under various traffic scenarios

We investigate the power performance of both the scheduled PSM and unscheduled PSM in this subsection. To better study the effectiveness of the scheduled PSM, a concept of optimal PSM is introduced. The optimal PSM has the theoretical optimal power performance in the sense that it only wakes up the mobile station in periods for receiving data from AP, and let the mobile station stay in sleep state in remaining periods. The ACK is reserved but the RTS/CTS are eliminated for the optimal PSM. We compare the power consumption of the three PSMs under various traffic scenarios.

As a start, we examine the power performance of the three PSMs when no background traffic is imposed on the network. The foreground traffic is a downlink stream with its data rate adjusted from 0kbps to 4000kbps. Both multicast and unicast cases are considered for the foreground stream. Fig. 5 shows the power consumption of the wireless interface as a function of the data rate of the foreground traffic. In all traffic cases, the unscheduled PSM consumes more power than scheduled PSM, and the gap becomes larger when data rate increases, particularly for unicast foreground traffic. As shown in Fig. 5(a), compared to unscheduled PSM, the net saved power by scheduled PSM reaches as much as 0.4W at the rate of 3000kbps, accounting for almost 50 percent of the total power consumption by unscheduled PSM, whereas in Fig. 5(b) this value is reduced to only 0.08W. This difference lies in that more efforts are made for unicast transmissions to reduce power consumption in scheduled PSM than for multicast transmissions. In multicast, the saved power mainly comes from the shortened time of medium access by adoption of bursty frame delivery in a SP. However, in unicast, in addition to that, the power consumption is further reduced by using data transmission prefix RTS/CTS for each service period instead of each frame, which saves energy for frequent RTS/CTS exchange. It can be inferred that the latter effect on the power consumption is much greater than the former. Moreover, in both unicast and multicast cases, the performance of scheduled

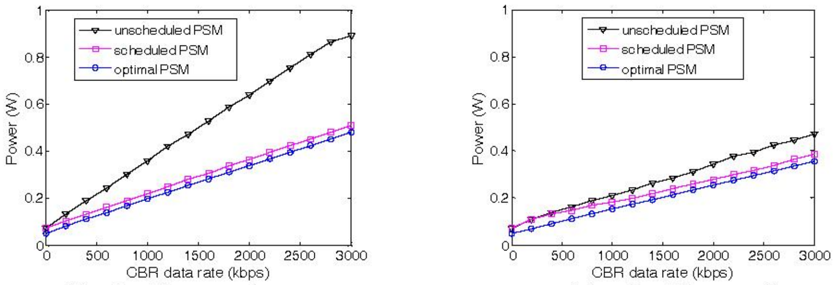

(a) unicast(foreground) (b) multicast(foreground)

Fig. 5 Power consumption as a function of data rate of the foreground traffic.

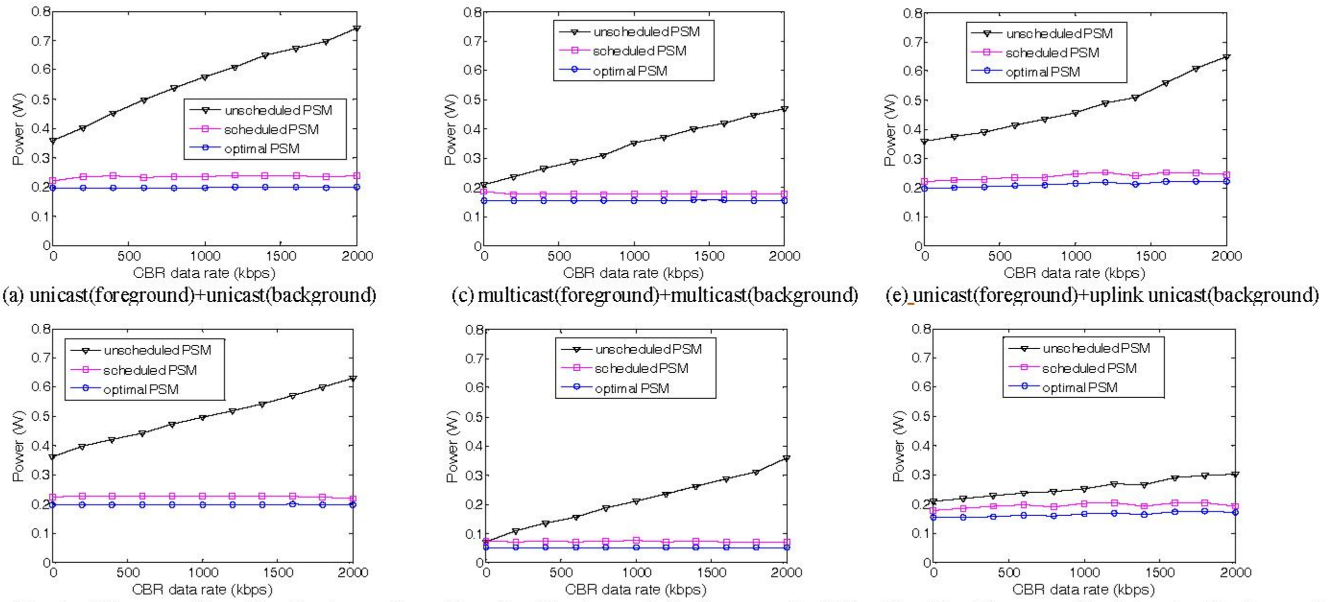

(a) unicast(foreground)+unicast(background) (c) multicast(foreground)+multicast(background) (e) unicast(foreground)+uplink unicast(background)

(b) unicast(foreground)+multicast(background) (d) multicast(background) (no foreground traffic) (f)multicast(foreground)+uplink unicast(background)

Fig. 6 Power consumption as a function of data rate of the background traffic.

PSM closely approaches to that of optimal PSM, with an average gap no more than 0.025W. Hence the scheduled PSM behaviors in a highly energy-efficient manner for traffic scenarios without background traffic.

To study the power performance of these PSMs in circumstances with background traffic, we create two streams connecting AP with station 2 and 3 separately. For each stream, traffic data rate is tuned to increase from 0 to 1000kbps, thus the total data rate of background traffic varies from 0 to 2000kbps. Multiple traffic types are considered in our simulations to reflect the impact of various background traffics, including downlink unicast, downlink multicast and uplink unicast. When multiple multicast streams coexist in the network, they are supposed to belong to different multicast groups. The foreground traffic is a downlink stream sourced at AP and destined to station 1, with a fixed data rate at 1000kbps except in one case (Fig. 6(c)). Two traffic types are considered for foreground traffic, downlink unicast and downlink multicast. We show the power consumption as a function of data rate of the background traffic in various scenarios in Fig. 6.

Fig. 6(a) and Fig. 6(b) investigate the impact of background traffics on the power performance of station 1 running a downlink unicast stream. In Fig. 6(a), the background traffic is two downlink unicast streams, while in Fig. 6(b), the background traffic is two downlink multicast streams. In both two cases, the scheduled PSM defeats the unscheduled PSM in terms of power conservation. For scheduled PSM, the average power consumption maintains on a quite low level and keeps relatively steady across in all cases. The imposed background traffic, regardless of its type as unicast or multicast, causes no evident fluctuation of the power consumption. In contrast, for unscheduled PSM, the upward trend of the curves in both two graphs demonstrates that its power performance is susceptible to be affected by the background traffic. The impact of the unicast background traffic is even slightly more evident than that of multicast. Note that there is a gap between two PSMs when no background traffic is imposed, because the scheduled PSM saves more power than unscheduled PSM for foreground unicast traffic, which is demonstrated in Fig. 5.

Such a distinguishing performance of the two PSMs results from the difference of the handling procedure against the buffered frames. The unscheduled PSM treats each pending frame as an individual unit in the delivery process, and the time of the occurrence of a delivery can not be anticipated. While the former degrades its efficiency, the latter makes the mobile station switch the state of network interface in a passive manner. Both two factors make the mobile stations remain awake for a longer time than necessary and be susceptible to by affected by the network activities of other stations. The scheduled PSM handles the buffered frames in a more intelligent fashion. It groups these frames based on their destinations, and delivers them in batch at appointed time scheduled by the AP. Thus it is possible for the mobile stations to change the state of the network interface voluntarily based on the known arrangements of frame delivery. Further, the total waking up time is reduced to such a low level as it only accommodates frame delivery relevant to itself. During the period of irrelevant frame deliveries, it stays in a sleep state. Therefore, the power consumption for scheduled PSM is much less than that for unscheduled PSM. It does not fluctuate with the data rate of background traffics.

The conclusion is further confirmed in another scenario where both the foreground and background traffics are multicast streams. Fig. 6(c) shows the results. As can be seen in the graph, the power consumption for unscheduled PSM linearly increases with the imposed multicast traffic rate, while that for scheduled PSM remains quite steady throughout all data rates. The gap between the two reaches up to 0.292W at data rate of 2000kbps, whereas the power consumption for scheduled PSM is merely 0.176W at that point. Thus the reduced energy by scheduled PSM is considerable. Compared Fig. 6(b) with Fig. 6(c), the impact of background multicast traffic on unicast is more remarkable than that on multicast, and there is a shift between them. This shift may be explained by results in Fig. 5 that the gap of power consumption between the two PSMs for unicast is much bigger than that for multicast.

Note that in figure Fig. 6(c), though each multicast stream belongs to different multicast group, station 1 has to stand to listen to all multicast frames in the air for the unscheduled PSM. And even if the targeting station does not join any multicast group, it is required to keep awake during the whole multicast delivery procedure. We investigate this case in Fig. 6(d). As the data rate of the background traffic grows from 0 to 2000kbps, the power consumption of the station 1 running unscheduled PSM increases linearly from 0.07W to 0.36W, doubling 2 times for such a moderate background traffic case. And our further investigation shows that during this procedure, the percentage of sleeping time declines from nearly 100% to about 33%, implying that about 2/3 of deserved sleeping time is deprived by the imposition of the background traffic. For the scheduled PSM, both the power consumption and percentage of sleeping time maintains on a steady level in all data rate cases. Such a contrast shows the necessity and benefits for a power saving scheme to differentiate multicast groups.

In Fig. 6(e) and Fig. 6(f), we investigate the impact of the background uplink traffic on a mobile station's power consumption. The foreground traffic is a downlink unicast stream in Fig. 6(e) and a downlink multicast stream in Fig. 6(f). In both two cases, the scheduled PSM out-performs the unscheduled PSM in terms of power conservation and resistance to the impact of the background traffic, particularly when the foreground traffic is unicast. The increased power consumption for the unscheduled PSM is mainly contributed by the interruptions of transmissions from the uplink traffics, whereas the scheduled PSM does not have such a problem.

In all traffic scenarios, the scheduled PSM exhibits a close approach to the optimal PSM. The gap of power consumption between the two is bounded to 0.025W for all shown cases in Fig. 6. This result demonstrates that the scheduled PSM has a near optimal power performance.

C. *Network performance in terms of one-way delay*

Reducing the power consumption often comes with the cost of degradation of network performance. To simulate the network performance of the scheduled PSM and the unscheduled PSM, we set up three streams with fixed data rate at 200kbps connecting the AP and station 1 serving as foreground traffic, including a downlink unicast stream, a downlink multicast stream and a uplink unicast stream. In addition, two downlink streams are imposed between AP and station 2 and 3 serving as the background traffic. Two cases are considered for the background traffic, unicast and multicast, and they are examined in parallel for comparison in our simulations. We vary the data rate of each background stream from 0 to 1500kbps, and measure the network performance on each of the three foreground streams.

Among many metrics of network performance, we choose one-way delay (OWD) as the metric used for evaluation of network performance, because the negative effect of a power saving solution on network performance mainly results in the traffic delay between AP and mobile stations. The delay is caused by the buffering of relaying frames at the AP for the mobile stations in sleep state. A longer delay usually means more power conservation. A good power saving solution usually makes an appropriate tradeoff between the two. As for throughput, the scheduled PSM scheme also offers a slight improvement over the un-scheduled PSM due to the absence of Poll-Response sequence.

Fig. 7 shows the measured OWD of the foreground traffic as a function of the data rate of background traffics. Each column deals with one foreground stream: Fig. 7(a) and Fig. 7(b) deal with the downlink multicast stream; Fig. 7(c) and Fig. 7(d) deal with the downlink unicast stream; and Fig. 7(e) and Fig. 7(f) deal with the uplink unicast stream. In each column, the upper figure shows the OWD when the background traffic is unicast, and the lower figure shows the OWD when the background traffic is multicast. Results show that, in all situations, the OWD for scheduled PSM exceeds that for unscheduled PSM. In the first four figures, the gap between the two PSMs fluctuates slightly with the data rate of background traffic, and it is bounded to 15ms, which is approximately *4.1%-20.3%* of the OWD for unscheduled PSM. However, in the last two figures, the OWD for scheduled PSM increases remarkably with data rate of background traffics particularly for the case of multicast background traffic, while for unscheduled PSM it increases slowly. The gap between the two PSMs reaches up to *4.6ms* for unicast background traffic and *8.8ms* for multicast background traffic at a data rate of 3000kbps, approximately

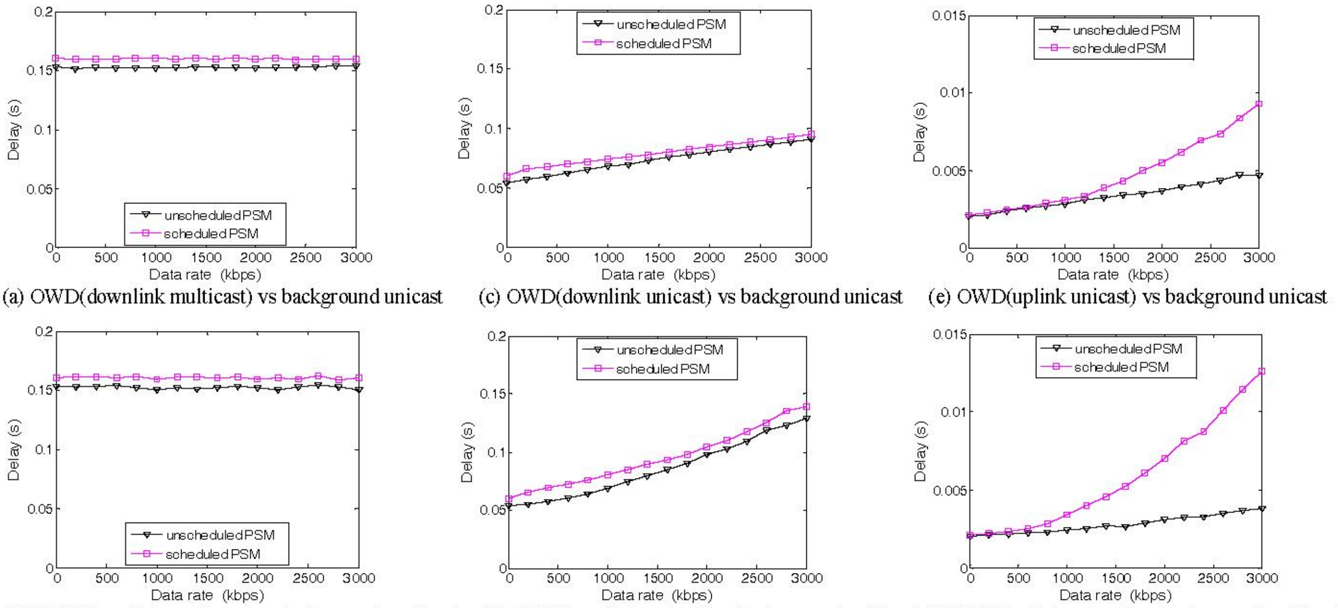

(a) OWD(downlink multicast) vs background unicast (c) OWD(downlink unicast) vs background unicast (e) OWD(uplink unicast) vs background unicast

(b) OWD(downlink multicast) vs background multicast (d) OWD(downlink unicast) vs background multicast (f) OWD(uplink unicast) vs background multicast

Fig. 7 One-way delay of the foreground stream as a function of the data rate of the background traffic.

97.9% and *231.2%* of the OWD separately for the unscheduled PSM at that point. These results demonstrate that the scheduled PSM has a more significant impact on the network performance of the uplink traffics than downlink. This difference is due to the high priority placed at the AP by the scheduled PSM over other mobile stations to access the medium. The AP can easily gain control of the medium and use it for a long time while the mobile stations have to delay their uplink frame transmissions. However, even so, the one-way delay of uplink streams for light to moderate background traffic is still less than 15ms, an acceptable level for most applications. If we adjust the value of surplus bandwidth allowance (A) in (1) to a large value, then the mobile stations would get more opportunities for uplink transmissions and the uplink OWD would be reduced. In general, with a given A, the downlink OWD gap will be bounded, and the uplink OWD gap will increase remarkably with the increase of traffic rate.

V. RELATED WORK

To the best of our knowledge, our scheduled PSM is the first power saving protocol that aims at addressing the issue of the degradation of power performance caused by background traffics. It costs only a small increase of beacon overhead but achieves near optimal power performance as demonstrated in our simulation results in section IV.

Subsequent to the unscheduled PSM, multiple power saving protocols have been proposed in the family of 802.11 standards. In 802.11e [2], an automatic power-save delivery (APSD) protocol is introduced. It defines two modes based on the delivery mechanism, namely the unscheduled APSD (U-APSD) and the scheduled APSD (S-APSD). In U-APSD, if a mobile station is informed that it has data outstanding, it issues a trigger frame to the AP and the AP responds with the frames buffered for it. In S-APSD, the unicast pending data is scheduled periodically by the AP to deliver at designated time, given that the mobile station has negotiated with the AP about the scheduling intervals. Both modes aim at improving the energy utilization via bursty frame transmissions. However, the un-coordinated trigger frames can cause the mobile station to wait for excessively long time before the data transmission commences. Meanwhile, the S-APSD relies on the resource reservation per stream by negotiation between AP and the mobile station, which degrades the flexibility compared to U-APSD and is not fit for the traffics with variable bit rate. Further, both of them don't address the issue of multicast data.

In the 802.11n draft [3], the APSD is enhanced by a power save multi-poll (PSMP) procedure. Similar to APSD, the PSMP procedure can be triggered by a frame (U-PSMP) or scheduled periodically by the AP (S-PSMP). It starts with broadcasting a PSMP message by the AP to the network which contains the scheduling information for downlink and uplink transmissions. The mobile station is assumed to wake up to receive this message and arranges its network activities and power state transitions based on the scheduling information. The PSMP creates effective transmission periods for broadcast, multicast and unicast data, and power consumption is greatly reduced by periodical occurrence of the PSMP procedure and its scheduled frame transmissions. However, compared to the scheduled PSM, the overhead of the management frames is considerable and the mechanism to maintain the uplink transmissions is fairly complex. Moreover, because the triggering mechanism can not guarantee that all members of the B/M group remain awake so as not to miss the data, the U-PSMP actually does not provide sufficient support for B/M services. Meanwhile, the S-PSMP deals with traffic on a per-stream level, which does not fit for the bursty traffic such as web browsing.

The 802.11v draft [4] proposes a flexible broadcast/multicast service (FBMS) to provide flexible management over the B/M traffics in power saving protocols. In FBMS, a B/M stream can choose to be served at a configurable delivery interval, which is multiple of DTIM periods. Thus traffics of different B/M groups can be differentiated by selecting different delivery intervals. The mobile station does not have to pay extra energy to receive irrelevant B/M traffics. The major disadvantage of the FBMS is that it significantly increases the delay of the B/M streams, and the saved energy is reduced when the number of B/M streams increases, as in this case some streams have to be delivered within the same DTIM periods. It does address the regular unicast traffic.

In addition to these standards, many researchers have explored power management strategies in wireless LANs. Simunic *et al.* [6] propose a time-indexed semi-Markov decision process model (TISMDP) to derive an optimal policy for system-level power management in portable systems. In [7] and [8], the authors investigate application-specific protocols to reduce the power consumption of network interfaces for streaming media applications. Kravets and Krishnan [9] present the design and implementation of a power-aware transport level protocol to achieve power conservation. In [10], the authors address the issue of interaction between energy-saving protocols and TCP performance for web-like transfers, and present a bounded slowdown (BSD) protocol that guarantees bounded delay on transfer round-trip time while conserves energy. This work is extended in [11], where a smart PSM is proposed to guarantee arbitrary user-desired delay performance. In [13], the network performance in terms of response time of HTTP applications is improved by using different beacon periods for different HTTP clients. In [12], the authors implemented a self-tuning power management (STPM) module in the linux kernel that adaptively adjusts the power management schemes with application hints. In [14], a survey of power saving protocols is provided. Though they have done excellent work to improve power or network performance in wireless LANs, none of them addresses the vulnerability of 802.11 PSM against the background traffic.

VI. Summary and Conclusion

Power conservation is a general concern for mobile computing. In this paper, we investigate the power performance of the legacy unscheduled PSM in an environment of 802.11 wireless LANs through analysis, experiments and simulations. We find that a mobile station's power performance is susceptible to the impact of the background traffics when multiple traffics are imposed on the network. Its power consumption grows substantially with the data rate of the background traffic. Even in case the mobile station does not join any multicast group, it consumes extra power for receiving the multicast data that are buffered at the AP for power saving. This vulnerability prevents the legacy unscheduled PSM from achieving high energy-efficient performance in networking scenarios with multiple traffic streams.

Considering the deficiencies of the legacy unscheduled PSM, we have designed a novel power saving protocol with scheduled, uninterrupted frame delivery procedure and differentiated multicasting in this paper, namely scheduled PSM. In the scheduled PSM, a beacon period is divided into multiple time slices, and all pending frames are scheduled to deliver at the appointed time slices. The TIM structure is redesigned to convey the slice assignment information. A mobile station can learn from the TIM element about the starting and ending time of the expected service period, expressed in time slices, and thus it is possible for the mobile station to control the state of the wireless interface adaptively to the traffic arrangements. To evaluate the performance of the scheduled PSM, a power saving module is implemented on ns2 and comprehensive simulations have been conducted. The results show that the scheduled PSM provides significant energy conservation over the unscheduled PSM particularly when multiple traffic streams coexist in the network, and it achieves this at cost of only slightly degradation of one-way delay performance.

References

[1] IEEE Computer Society LAN MAN Standards Committee. IEEE Std 802.11: Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications, August 1999.

[2] IEEE Computer Society LAN MAN Standards Committee. IEEE Standard 802.11e: Amendment to IEEE Std. 802.11: Medium Access Control (MAC) Quality of Service Enhancements, November, 2005.

[3] IEEE Computer Society LAN MAN Standards Committee. IEEE Standard 802.11n: Amendment to IEEE Std. 802.11: Enhancements for higher throughput, draft 1.06, 2006.

[4] IEEE Computer Society LAN MAN Standards Committee. IEEE Standard 802.11v: Amendment to IEEE Std. 802.11: Wireless Network Management, draft 0.05, 2006.

[5] L. M. Feeney and M. Nilsson, "Investigating the energy consumption of a wireless network interface in an Ad hoc networking environment," in Proc. Of IEEE Infocom, Anchorage, AK, April 2001.

[6] T. Simunic, L. Benini, P.W. Glynn, and G. D. Micheli, "Dynamic power management for portable systems," In Proc. ACM Mobicom, pages 11-19, Boston, MA, 2000.

[7] S. Chandra and A. Vahdat, "Application-specific network management for energy-aware streaming of popular multimedia format," In Proc. of the USENIX Annual Technical Conference, June 2002.

[8] S. Chandra, "Wireless network interface energy consumption implications of popular streaming formats," In Proc. Multimedia Computing and Networking (MMCN), pages 85-99, San Jose, CA, January 2002.

[9] R. Kravets and P. Krishnan, "Application-Driven Power Management for Mobile Communication," ACM Wireless Networks, vol. 6, no. 4, pp.263–277, 2000.

[10] R. Krashinsky, H. Balakrishnan, "Minimizing Energy for Wireless Web Access with Bounded Slowdown," In Proc. ACM Mobicom 02, pp. 119-130, Atlanta, GA, 2002.

[11] D. Qiao, K. G. Shin, "Smart power-saving mode for IEEE 802.11 wireless LANs," In Proc. Infocom, Miami, FL, 2005.

[12] M. Anand, E. B. Nightingale, and J. Flinn, "Self-tuning wireless network power management," Proc. ACM Mobicom 03, pages 176--189, San Diego, CA, September 2003.

[13] Z. Anderson, S. Nath, and S. Seshan, "Chooseing beacon period for improved response time for wireless http clients," in Proc. ACM MobiWac, Philadelphia, PA, October 2004.

[14] C. E. Jones, K. M. Sivalingam, P. Agrawal, and J. C. Chen, "A survey of energy efficient network protocols for wireless networks," ACM Wireless Networks, 7(4):343-358, July 2001.

[15] Ixchariot. http://www.ixiacom.com/products/display?skey=ixchariot.

[16] "The Network Simulator – ns-2," http://www.isi.edu/nsnam/ns/, online link.