

Hierarchical Power Management in Disruption Tolerant Networks with Traffic-Aware Optimization

Hyewon Jun Mostafa H. Ammar Mark D. Corner† Ellen W. Zegura

College of Computing, Georgia Institute of Technology, Atlanta, GA
{hyewon, ammar, ewz}@cc.gatech.edu

†Department of Computer Science, University of Massachusetts, Amherst, MA
mcorner@cs.umass.edu

ABSTRACT

Recent efforts in Disruption Tolerant Networks (DTNs) have shown that mobility can be a powerful means for delivering messages in highly-challenged environments. DTNs are wireless mobile networks that are particularly useful in sparse environments where the density of nodes is insufficient to support direct end-to-end communication. Unfortunately, many mobility scenarios depend on untethered devices with limited energy supplies. Without careful management depleted energy supplies will degrade network connectivity and counteract the robustness gained by mobility. A primary concern is the energy consumed by wireless communication, and in particular the energy consumed in searching for other nodes to communicate with. In this architecture, energy can be conserved by using the low-power radio to discover communication opportunities with other nodes and waking the high-power radio to undertake the data transmission. We develop a generalized power management framework for controlling the wakeup intervals of the two radios. In addition, we show how to incorporate knowledge of the traffic load, and we devise approximation algorithms to control the sleep/wake-up cycling to provide maximum energy conservation while discovering enough communication opportunities to handle that load. We evaluate our schemes through simulation and compare them against single radio architectures, and against algorithms that do not incorporate information about the load. Our results show that the generalized power management mechanism achieves better energy efficiency in discovering communication opportunities than mechanisms relying on only one radio. Also, when traffic load can be predicted, our approximation algorithm could reduce energy consumption from 73% to 93% compared with the case without power management. Finally, our results show that the statistical information about traffic load and network connectivity allows us to devise an efficient algorithm even for the one-radio architecture, so energy can be saved almost equivalently to the two radio architecture.

1. INTRODUCTION

Work supported in part by NSF Grants NETS-0519784, CNS-0519881, CNS-0520729, CNS-0447877, and ITR-0313062.

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.

SIGCOMM'06 Workshops September 11-15, 2006, Pisa, Italy.
Copyright 2006 ACM 1-59593-417-0/06/0009 ...\$5.00.

Network designers often think of mobility as a detriment to building robust networks. However, recent efforts in Disruption Tolerant Networks (DTNs) have shown that mobility can be a powerful means for delivering messages in highly-challenged environments [10, 14, 23, 13, 32, 16, 6]. DTNs are wireless mobile networks that are particularly useful in sparse environments where the density of nodes is insufficient to support direct end-to-end communication. When mobile nodes encounter each other, opportunistically or intentionally, they pass messages to route them toward their final destinations. Unfortunately, many mobility scenarios depend on untethered devices with limited energy supplies. Without careful management depleted energy supplies will degrade network connectivity and counteract the robustness gained through mobility.

However, mobile devices exhibit a tension between saving energy and providing connectivity through opportunistic encounters. In order to pass messages, the device must discover *contacts* with other nodes [14]—typically the discovery is done using the same wireless interface used for message transfer. At the same time, energy conservation requires disabling (i.e., sleeping) the wireless device—the wireless interface is one of the largest energy consumers in mobile devices [20] whether they are actively communicating or just listening [11, 28]. If the wireless interface is asleep, the node cannot discover other nodes for communication and vice versa. Thus, power management in DTNs must balance discovery of other nodes while aggressively disabling the radio during the remaining periods.

In a previous study, we have shown how nodes can effectively use information about network connectivity to predict when to enable their wireless interfaces and search for contacts [18]. However, those results only applied to a network with a certain degree of regularity, for instance a node that followed a known schedule. Thus, a network with significant randomness in node mobility requires different mechanisms to save energy while delivering messages.

In this paper, we examine the possibility of using a hierarchical radio architecture in mobile DTNs, in which nodes are equipped with two complementary radios: a long-range, high-power radio and a short-range, low-power radio. In this architecture, energy can be conserved by using the low-power radio to discover contacts with other nodes and then waking the high-power radio to undertake the data transmission. Previous studies using this hierarchical radio architecture have considered only densely deployed networks, in which the short range of the low power radio is sufficient to discover each other [24, 27, 21]. However, DTNs are generally applicable to sparser networks, where the low-power radio may reach a subset of other nodes, even when the nodes are mobile. Therefore, if a node relies only on the low power radio to

discover contacts, it may miss them due to the shorter range. To avoid missing contacts, we propose a generalized power management scheme that uses *both* radios to participate in contact discovery. This generalized scheme controls the *wake-up interval* of each radio, which it uses to trade between energy savings and the performance of message delivery. To guide this trade-off, we also propose incorporating knowledge of the traffic load in the network. We devise approximation algorithms to determine the optimal wake-up intervals that minimize the overall energy consumption, while discovering enough contacts to handle the expected traffic load.

We compare our generalized scheme with two alternative schemes: one that only uses the high power radio for discovery, and one that only uses the low power radio. In Section 3 we show that without traffic information, the generalized two-radio scheme, achieves equivalent, or better energy efficiency when compared to single radio designs. However, these gains depend heavily on what the wakeup interval of the radio is set to. In Section 5 we show that by using the traffic load to inform the wakeup interval our approximation algorithms reduce energy consumption by 73% to 93% compared to the case without power management. Finally, our traffic-aware algorithm for the one-radio architecture can save almost equivalent energy to those for the two-radio architecture, lowering the overall benefits of the two-radio architecture.

While this seems to contradict the results of previous studies [24, 27, 21] and our results without traffic load information, the difference lies in the use of additional knowledge. This knowledge about the traffic load allows us to devise efficient algorithms even for the one-radio architecture to overcome the disadvantage of not having the additional radio. In sum, if the system does not have knowledge about the traffic load, then two radios can be used to save energy. However *with* such knowledge the system can optimize a single radio design to discover contacts at efficiencies similar to two radio designs.

The remainder of this paper is organized as follows. In Section 2, we presents our power management framework. Section 3 shows our simulation results to evaluate the performance of our power management under various wake-up intervals. In Section 4, we propose wake-up interval estimation for our power management according to the expected traffic load. Section 5 illustrates our simulation results to evaluate the traffic-aware wake-up interval estimation. We discuss related work in Section 6 and conclude this paper in Section 7.

2. DTN RADIO DISCOVERY

In a mobile DTN, two nodes communicate with one another during *contacts* [14] that occur when the two nodes, either mobile or stationary, are within the radio range of one another. If the devices are equipped with multiple radios, the nodes may or may not discover a particular contact opportunity depending on the range of the radios, which radios are active, and the movement trajectory of the two nodes. Figure 1 shows two possible contact scenarios. In Figure 1(a), node A moves along the trajectory shown and node B stays in one location. Node A first enters within B’s high-power radio range and then within B’s low-power radio range. As a result, the nodes have a long contact via their high-power radios and a short contact via their low-power radios. In contrast, Figure 1(b) shows that when node A passes node B using a different trajectory, it only enters the range of the high-power radio of node B. As a result, the nodes have only one contact via their high-power radios.

While the short range radio may discover less contact opportunities than the long range radio, it does so at substantially reduced energy costs. Table 1 shows the energy consumption of two sample radios, one high-power radio: 802.11, and one low-power ra-

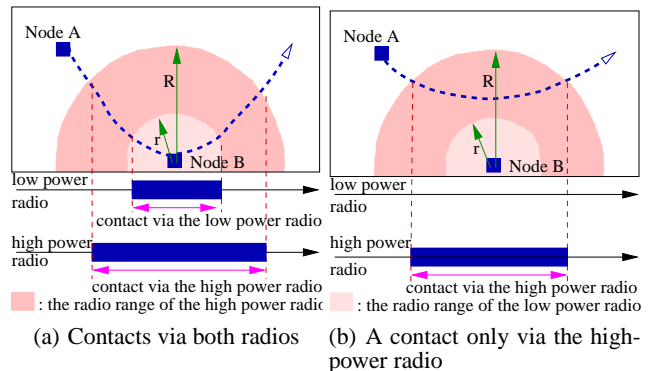


Figure 1: Contacts discovered by the high-power and low-power radios, where R and r are the radio radii of the high-power and low-power radios, respectively

Table 1: Power usage of a Lucent IEEE 802.11 WaveLAN PC Card and a Chipcon CC1000 (unit:W)

Activity	Transmit	Receive	Idle	Sleep
WaveLan PC card	1.3272	0.9670	0.8437	0.0664
Chipcon CC1000	0.0781	0.0222	0.0222	0.00003

dio: the CC1000 radio found in Mica2 Motes [11, 8, 21]. The energy consumption of each wireless interface depends on whether it is *transmitting*, *receiving*, *idling* (when listening to the wireless medium without transmitting nor receiving), or *sleeping* (when the wireless interface is disabled). When sleeping, a node consumes an order of magnitude less energy than when idling, while an idling node consumes energy at the same order of magnitude as a receiving or transmitting node. In addition, the CC1000 radio consumes an order of magnitude less energy than the 802.11 radio for each activity. Thus, it can discover contacts using substantially less power. However, its outdoor range is limited to 50-100m, while the 802.11 radio has a range of 250-500m.

In this paper, we consider a DTN consisting of mobile nodes as well as stationary nodes, which have two radio interfaces: one with a long radio range and high-power, e.g., a 802.11 wireless card, and the other with a short radio range and low-power, e.g., a CC1000 radio. We only account for the communication energy consumption of a wireless interface and do not consider other sources such as computation or mobility. Also, we assume that nodes have no a-priori knowledge about other nodes’ mobility and contacts must be discovered by one or both radios of the nodes. To discover contacts, a radio broadcasts messages called *beacons* periodically. To save energy while discovering contacts, a radio has three power management modes: *search*, *contact*, and *dormant modes*. In the search mode, the radio wakes up periodically to discover a contact. This period is called a *wake-up interval*. In the contact mode, the radio stays awake to exchange messages with other nodes that it previously discovered in the search mode. In the dormant mode, the radio is not used and remains asleep.

Given that the high-power radio is always used, there are four possible variations of this general framework. Table 2 summarizes the power management mechanisms used in this paper. The *Continuous Aware Mechanism (CAM)* uses only the high-power radio of a node. In CAM, the high power radio always stays awake to search for other nodes. The *Power Saving Mechanism (PSM)* also uses only the high-power radio, however it alternates between be-

ing asleep and being awake while discovering contacts. While similar, note that this PSM is not the same as 802.11's PSM mode. The *Short-range-radio-dependent Power Saving Mechanism (SPSM)* uses both low-power and high-power radios of a node. In this mechanism, the low-power radio alternates between sleeping and waking up to discover contacts, while the high-power radio sleeps and is awakened by the low-power radio only after discovering a contact. Finally, the *Generalized Power Saving Mechanism (GPSM)* uses both the low power and high-power radios of a node. In this mechanism, both radios alternate between sleeping and waking up to discover contacts. If a contact is discovered by the low-power radio, the high-power radio is awakened. If a contact is discovered by the high-power radio, the radio stays awake as long as it has contact with the other node. In the remainder of this section, we describe each of these mechanisms in detail, with the exception of the rather straightforward CAM.

2.1 Power Saving Mechanism (PSM)

In PSM, a node uses only one long-range radio and its radio periodically sleeps and wakes up to save energy while discovering contacts. Such a scheme was first described in our previous work [18] and we summarize it below for completeness.

We describe the wake-up behavior using three quantities: *beacon window*, *beacon period*, and *wake-up interval*. The beacon window is a time period when the radio wakes up and chooses a random time to transmit a beacon. A beacon period is the time between the beginning of two consecutive beacon windows in the contact mode. Finally, a wake-up interval is the time between two consecutive wake-ups in the search mode, which can be enlarged as a multiple of a beacon period. Note that we assume that nodes have synchronized clocks from a source such as GPS, and thus the beacon windows can be synchronized to start on common, discrete intervals.

These time periods are used as illustrated in Figure 2 when a node contacts only one node. Initially, the radio of a node is in the search mode since it is not in contact with node k . In the beginning of a beacon window, it chooses a random time within the beacon window to broadcast a beacon. If it does not receive another node's beacon, the radio sleeps at the end of the beacon window and wakes up at the beginning of the next beacon window. If the radio receives a beacon, it enters the contact mode, in which it stays awake to exchange messages. In the contact mode it continues to send a beacon at every beacon period and listens for beacons from the same node. If the radio fails to receive a certain number of beacons consecutively, it considers the contact is terminated and returns to the search mode. When a node contacts multiple nodes, its radio should consider multiple contact states to transit between the power management modes. The detailed procedure and algorithm can be found in [18].

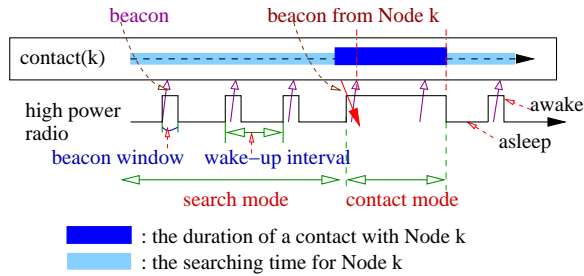


Figure 2: Transition between power management modes when a node has contacts with only one node, node k , using PSM

2.2 Short-range-radio-dependent Power Saving Mechanism (SPSM)

In SPSM, a node depends only on the low-power radio to discover contacts. As in PSM, nodes have synchronized clocks and both radios have their own three time periods to synchronize contact discovery procedure among nodes.

This scheme is identical to PSM except that the low-power radio conducts PSM's search mode as illustrated in Figure 3. If the low-power radio detects a contact, it wakes up the high-power radio, which then enters the contact mode. In the contact mode, the high-power radio stays awake to exchange messages, using the beacons to detect when the contact ends. At the same time, the high-power radio may detect other contacts by receiving beacons. Meanwhile the low-power radio continues to search for other contacts. When all the contacts are terminated, the high-power radio returns to the dormant mode and the low-power radio continues to search.

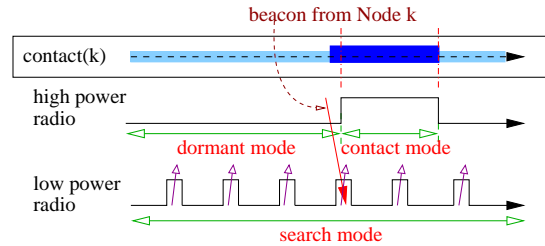


Figure 3: Transition between power management modes when a node has contacts with only one node, node k , using SPSM

2.3 Generalized Power Saving Mechanism (GPSM)

In GPSM, a node utilizes both radios to discover contacts. The two radios have separate wake-up intervals, beacon periods, and beacon windows. For instance, the high-power radio may have a larger wake-up interval than the low power radio, in which case the low-power radio searches for contacts more frequently than the high-power radio.

Figure 4 illustrates the power management with an example scenario. Initially, a node has no contact, so that both of its radios alternate between sleeping and waking up to send and listen for beacons in the search mode. If the low-power radio receives a beacon from node k , it wakes up the high-power radio and keeps searching for other contacts. If the high-power radio receives a beacon from node k , it enters the contact mode and stays awake to exchange messages. It also sends and listens for beacons from the same node. If the high-power radio does not receive a certain number of beacons consecutively from the same node, it considers that the contact is terminated and enters the search mode. When a node makes contact with multiple nodes, the high-power radio in the contact mode may detect other contacts by receiving beacons.

This mechanism is a generalized scheme that includes PSM and SPSM at its two extremes. If we set the wake-up interval of the low-power radio to infinity, GPSM works like PSM. Also, if we set the wake-up interval of the high-power radio to infinity, it works like SPSM.

Table 2: Power management mechanisms depending on how the radios are used

Power management	low-power radio	high-power radio
CAM	Not used	Always on
PSM	Not used	sleep/wake-up cycling to discover contacts
SPSM	Sleep/wake-up cycling to discover contacts	Awakened by the low-power radio
GPSM	Sleep/wake-up cycling to discover contacts	Sleep/wake-up cycling to discover contacts, <i>and</i> Awakened by the low-power radio

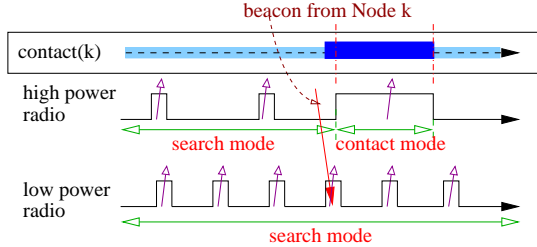


Figure 4: Transition between power management modes when a node has contacts with only one node, node k , using GPSM

3. PERFORMANCE COMPARISON

To provide some insight into the relative performance of these schemes, we provide a set of simulation results. Specifically, we investigate the impact of the wake-up interval of the high-power radio on the contact discovery performance. We consider two metrics: (1) *contact discovery ratio*, which is the discovered contact time divided by the total contact time that the high-power radio could have discovered when using CAM, and (2) *energy efficiency*, which is the average amount of discovered contact time per unit energy used. In other words, we measure how much energy the system spends to find a certain amount of message transfer time—the more transfer time discovered per joule of energy, the higher the efficiency is.

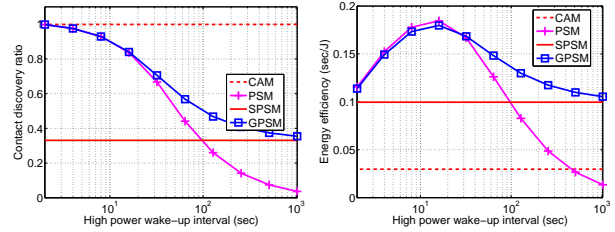
We use the ns-2 simulator with the following parameters. We use 802.11 MAC layer for both interfaces at different channels. Also, both interfaces set the beacon period to 2 seconds and the beacon window to 300ms. Each simulation runs for 50 hours of simulated time and each point of our graphs represents the average of ten runs. To simulate node movement, we use the Random-Waypoint (RWP) [15] mobility model. For the experiment we use the energy consumption of an 802.11 and a CC1000 radio as shown in Table 1. For the low-power radio, the radio range and the data rate are 100m and 76.8Kb/s, respectively. For the high-power radio, the radio range and the data rate are 250m and 2Mb/s, respectively.

In Figure 5, we show the impact of the wake-up interval of the high-power radio on the performance of contact discovery as the interval varies from 2 to 1024 seconds to show how it trades between energy savings and contact discovery performance. The wake-up interval of the low-power radio is set to 2 seconds for both SPSM and GPSM.

Figure 5(a) shows that the contact discovery ratio of PSM decreases from 100% to below 10% as the wake-up interval increases. On the other hand, GPSM discovers more than 37% of contacts for all wake-up intervals because the low-power radio assists the discovery. In addition, its discovery ratio is always more than that of SPSM because of the high-power radio.

Figure 5(b) shows that the energy efficiency of GPSM and PSM increases and then decreases as the wake-up interval decreases. Using smaller wake-up intervals than the peak of the graph wastes energy without discovering many more contacts. Also, Figure 5(b)

shows that GPSM is more energy efficient than PSM and SPSM for all wake-up intervals.



(a) contact discovery ratio (b) energy efficiency

Figure 5: The impact of wake-up intervals of the high-power radio in the RWP scenario

This experiment demonstrates that the wake-up interval of the radio greatly impacts energy efficiency, and the usefulness of the second radio is highly dependent on that interval. More importantly, this does not answer the question as to how to set the interval in the first place. In the next section, we provide this as an analytical result when additional information about contacts and traffic load is available.

4. TRAFFIC-AWARE WAKE-UP INTERVAL ESTIMATION

The critical issue in all of the power management mechanisms is determining the proper wake-up intervals. In each mechanism, the wake-up intervals can be used to trade between energy and delivery performance. Specifically, when the wake-up intervals are long, energy can be saved at the cost of missing contacts, which in turn results in poor delivery performance. When the wake-up intervals are short, more energy will be expended while discovering more contacts and improving delivery performance. The wake-up intervals can be increased for energy savings; however, the intervals cannot be too long because nodes may drop messages. If the generated traffic is greater than the delivery rate, the node will eventually drop packets due to its limited buffer space. To avoid dropping packets, we state the problem as follows: *For each node, find the wake-up intervals that leads to discovering enough contacts to deliver the traffic load while minimizing energy consumption.*

To formulate the problem, we assume that statistical information about contacts and traffic load between each pair of nodes is available. This information is often already available in a DTN: nodes observe and exchange the history of contacts and traffic load to make efficient routing decisions [6, 14].

To address the problem, we define the *contact arrival rate* as the number of contacts between two nodes over a unit time and *expected bandwidth* as the maximum amount of data that can be delivered by the discovered contacts between two nodes over a unit time. With these definitions, we can estimate the expected bandwidth for given wake-up intervals using contact duration and contact arrival rate per pair of nodes. Unfortunately, the distribution

Table 3: Notation used in the wake-up interval estimation

N	the set of nodes in a network
R	the radio range of the high-power interface
r	the radio range of the low-power interface
θ	the bandwidth of the high-power interface
f_{ij}	the traffic load on the link (i, j)
w^h, w^l	the wake-up intervals of the high-power and low-power radios, respectively

of inter-arrival times and contact durations is generally not known, especially for many common mobility models [4]. Without this, it is not possible to develop general, optimal algorithms. Therefore, we devise approximation algorithms assuming that contacts arrive according to a Poisson process and contact durations are constant and equal to the average contact duration from a mobility scenario. In Section 5, we validate this approximation and show that it works well in practice. The notation used in the wake-up interval estimation is summarized in Table 3.

4.1 Wake-Up Interval Estimation for PSM

Recall that PSM uses only the high-power radio to discover contacts. Since individual nodes have different traffic loads, we choose a different optimal wake-up interval for each node. However, given that nodes wake up at different rates, we must ensure that the nodes can still synchronize their wake-ups so that one node hears another node's beacon. To do this, we choose the wake-up intervals from a set that are a multiple of the interval between another node's beacons, which we refer to as the beacon period. Thus, we assume that a network operator provides a set of wake-up interval candidates as follows:

$$W^h = \{w^h \mid w^h = 2^k b^h, k = 0, 1, 2, \dots, K\}, \quad (1)$$

where b^h is a beacon period of the high-power radio and K is a non-negative integer.

From W^h , we determine the optimal wake-up interval of the high-power radio for node i as follows. First, we estimate the expected bandwidth on the link (i, j) for each wake-up interval $w^h \in W^h$, as illustrated in Figure 6. Initially, an opportunity to have a contact on the link (i, j) starts at time t after the termination of the last wake-up. We ignore a beacon window because it is very short. Because we assume that contacts arrive as a Poisson process, the arrival time of a contact is approximately uniformly distributed within a wake-up interval, i.e., $t \sim \text{uniform}(0, w^h)$. At the next wake-up, the contact can be discovered if the contact is longer than the remaining time until the next wake-up, i.e., $w^h - t < d_{ij}^h$, where d_{ij}^h is the contact duration on the link (i, j) that can be discovered by the high-power radio. In addition, the length of the discovered contact time is the contact duration less the remaining time until the next wake-up, i.e., $d_{ij}^h - (w^h - t)$. Thus, the expected discovered contact time $c_{ij}^{psm}(w^h)$ of a contact for a given w^h is calculated as follows. If $w^h \leq d_{ij}^h$,

$$c_{ij}^{psm}(w^h) = \int_0^{w^h} \frac{1}{w^h} (d_{ij}^h - w^h + t) dt. \quad (2)$$

Otherwise (i.e., if $w^h > d_{ij}^h$),

$$c_{ij}^{psm}(w^h) = \int_0^{w^h - d_{ij}^h} \frac{1}{w^h} \cdot 0 dt + \int_{w^h - d_{ij}^h}^{w^h} \frac{1}{w^h} (d_{ij}^h - w^h + t) dt. \quad (3)$$

The first integral corresponds to the case when the contact is shorter than the wake-up interval, and the contact opportunity does not overlap with any of the beacon windows. The second case corresponds to the case when the contact duration can be discovered by a wake-up, but the node misses the first part of the contact. As a result,

$$c_{ij}^{psm}(w^h) = \begin{cases} d_{ij}^h - \frac{1}{2}w^h & \text{if } w^h \leq d_{ij}^h \\ \frac{(d_{ij}^h)^2}{2w^h} & \text{otherwise} \end{cases} \quad (4)$$

Then, the expected bandwidth on the link (i, j) is $\lambda_{ij}^h \cdot \theta \cdot c_{ij}^{psm}(w^h)$ for a given w^h , where λ_{ij}^h is the contact arrival rate when the high-power radio is used on the link (i, j) , and θ is the bandwidth of the high-power radio.

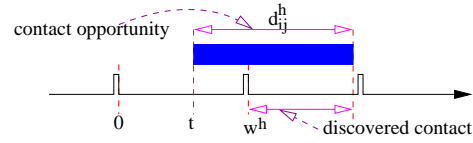


Figure 6: A discovered contact time by PSM for a given w^h

Secondly, we find the set of wake-up intervals, w^h , that provide bandwidth greater than or equal to the traffic load on the link (i, j) . We choose the wake-up interval from this set, by finding the interval that consumes the least amount of energy. In the case of PSM, this is straightforward as it is the maximum wake-up interval w_{ij}^h in the set, i.e.,

$$w_{ij}^h = \max \left\{ w^h \mid \lambda_{ij}^h \cdot \theta \cdot c_{ij}^{psm}(w^h) \geq f_{ij}, w^h \in W^h \right\}, \quad (5)$$

where f_{ij} is the traffic load on the link (i, j) . Then, the resulting w_{ij}^h is the wake-up interval that will consume the least amount of energy while satisfying the bandwidth constraint on the link (i, j) .

Since the node contacts multiple other nodes, and it does not know which node it will discover next, it must choose the minimum wake-up interval among all of the wake-up intervals it computed for each outgoing link, w_i^h , i.e.,

$$w_i^h = \min \{ w_{ij}^h \mid j \in N - \{i\} \}, \quad (6)$$

where N is the set of nodes in the network. Then, the resulting w_i^h is the wake-up interval that consumes the least amount of energy while satisfying the bandwidth constraints for all the links of node i . This result is optimal for the set of discrete wake-up intervals the algorithm considers.

4.2 Wake-Up Interval Estimation for SPSM

In contrast with PSM, SPSM relies solely on the low-power radio to discover contacts. As shown in Figure 1, the length of a contact via the low-power radio is shorter than that via the high-power radio. Thus, once a contact is discovered via the low-power radio, the high-power radio monitors the state of the contact and extends the contact time beyond the time the node leaves the range of the low-power radio. This extended amount of a contact is denoted as δ_{ij} for the link (i, j) . This δ_{ij} can be observed by having a

node maintain its contact states by both radios and estimate the time difference between the termination of their contacts, i.e., (the termination of a contact discovered by the high-power radio - that of the corresponding contact by the low-power radio). This extension makes the estimation of the expected bandwidth of SPSM different from that of PSM. Except for that, the wake-up interval estimation procedure of SPSM is identical to that of PSM. The detailed procedure of the wake-up interval estimation for SPSM can be found in our technical report [17].

4.3 Wake-Up Interval Estimation for GPSM

This power management uses both high-power and low-power radios to discover contacts. Since contacts can be discovered by either of the radios, the wake-up interval estimation procedure is more complicated than that of the other mechanisms. In brief, we determine the wake-up intervals for node i as follows. First, we estimate the expected bandwidth on the link (i, j) for each pair of wake-up intervals (w^h, w^l) in a given set of discrete wake-up interval pairs: i.e., $\{(w^h, w^l) \mid w^h \in W^h \cup \{\infty\}, w^l \in W^l \cup \{\infty\}\}$, where W^h and W^l are sets of finite discrete wake-up interval candidates provided by a network operator for the high power and low power radios, respectively, as in Section 4.1. Secondly, we find a set S_{ij} of wake-up interval pairs that provide enough bandwidth to consume the expected traffic load on the link (i, j) for each $j \in N - \{i\}$. Third, we find the pair of wake-up intervals (w_{ij}^h, w_{ij}^l) in S_{ij} that will consume the least amount of energy assuming that node i contacts only node j , so both nodes choose the same pair of wake-up intervals that satisfies the bandwidth requirement on their link and possibly consumes the least amount of energy. Finally, the node chooses the minimum wake-up intervals w_i^h and w_i^l among w_{ij}^h and w_{ij}^l , respectively. Then, the resulting (w_i^h, w_i^l) is the pair of wake-up intervals that consumes approximately the least amount of energy while satisfying the bandwidth constraints for all the links of node i . The detailed procedure can be found in our technical report [17].

5. PERFORMANCE EVALUATION

Our goal in evaluating our traffic-aware optimization process is to show two things. First, we show that our analytical model can accurately predict the amount of discovered contact time for any particular wake-up intervals. If the predicted contact time is correct, then the wake up intervals can be tuned to a particular traffic load. Second, we show that these wake-up mechanisms save significant energy in a DTN node and compare the use of single and multiple radio search mechanisms.

To show this we use the same Random Way Point model from earlier as well as a Message Ferrying (MF) mobility model [32]. In MF the network consists of 12 stationary nodes and eight mobile nodes, called ferries. We place stationary nodes in a grid of 5km \times 5km area sparsely. Then ferries visit each node in order and repeat the route, starting at different locations. The key difference between the RWP and MF scenarios is the distance that nodes pass one another: in the RWP scenario, contacts occur at much longer distances than in the MF scenario.

5.1 Contact Discovery

Figure 7 shows how close our estimation of contact discovery in Section 4 is in comparison to the actual contact discovery. The contact discovery ratio is the discovered contact time divided by the possible contact time that could have been discovered by using the high-power radio in an always-on mode (CAM). Each line represents the estimated contact discovery ratio from our model and the simulated contact discovery ratio of each power management in RWP

and MF scenarios as indicated. Each scenario has a pair of nodes in a deployment area and each simulation runs for 500 simulation hours. As Figure 7 shows, our estimation could predict the contact discovery ratio close to the simulation results. GPSM tends to have more error than other mechanisms due to its complexity, while its estimation is still close to the simulated results. These graphs are representative for other pairs of wake-up intervals in our evaluation.

5.2 Traffic-Aware Power Management

Given the accuracy in determining the contacts resulting from a particular wake-up interval, we compare the schemes using our traffic-aware wake-up interval estimation. We consider the following metrics: (1) *normalized energy consumption* that is the total energy consumption of each power management case divided by that of CAM, and (2) *delivery ratio* that is the ratio of successfully delivered messages to the total number of generated messages. A message loss occurs when a buffer overflows at a node in the middle of a routing path.

In simulations, a message loss occurs when a message buffer overflows at a node in the middle of a routing path because the node attempts to store messages beyond the limitation of its buffers. The routing path selection and traffic generation are described in our technical report [17].

Figure 8 shows the impact of traffic-aware optimization on the power management performance in the RWP and MF scenarios as the traffic load varies from 1.6Kb/s to 54Kb/s. We compare all of the schemes with CAM to observe how much energy can be saved. Also, we compare it with a PSM when the traffic load cannot be predicted. In such a case, we use a fixed wake-up interval that uses the least amount of energy to discover each contact. This allows us to compare a traffic-aware mechanism with one that neglects any knowledge of traffic and focuses solely on energy efficient discovery. In our scenarios, such a wake-up interval for PSM is 16 seconds and the corresponding curve is indicated by ‘‘PSM-16’’ in Figure 8.

In Figures 8(a) and (b), we illustrate the simulation results in the RWP scenario. Figure 8(a) shows that all mechanisms consume energy only from 7% to 27% compared to no power management. Also, it shows that as the traffic load becomes heavier, all mechanisms with traffic-aware optimization consumes more energy because they choose smaller wake-up intervals and end up discovering more contacts. Among these mechanisms, SPSM energy consumption is relatively flat. As the load grows, SPSM searches for contacts as aggressively as it can. However, SPSM relies on the low-power radio for contact discovery, which has insufficient range to discover enough contacts to handle its traffic. As a result, SPSM discovers less contacts, consuming less energy than other mechanisms as shown in Figure 8(a), and has a low delivery ratio when the traffic load is more than 5.3Kb/s as shown in Figure 8(b). On the other hand, PSM and GPSM achieve more than 90% delivery ratio for all traffic loads, which is close to the delivery ratio of CAM. In this RWP scenario, GPSM works similarly to PSM because nodes pass by each other at long distances and most of their contacts are discovered by the high-power radio. At the same time, GPSM saves energy up to 8.7% of the PSM case because its low-power radio also discovers contacts with less energy consumption. Finally, the energy consumption of PSM-16 is almost constant because of its fixed wake-up interval. The comparison with PSM-16 shows that power management with traffic-aware optimization saves more energy than PSM-16 while delivering the same amount of messages when the traffic load is light. When traffic load is heavy, PSM and GPSM consume more energy, but at the same time deliver more messages than PSM-16. Thus, the additional informa-

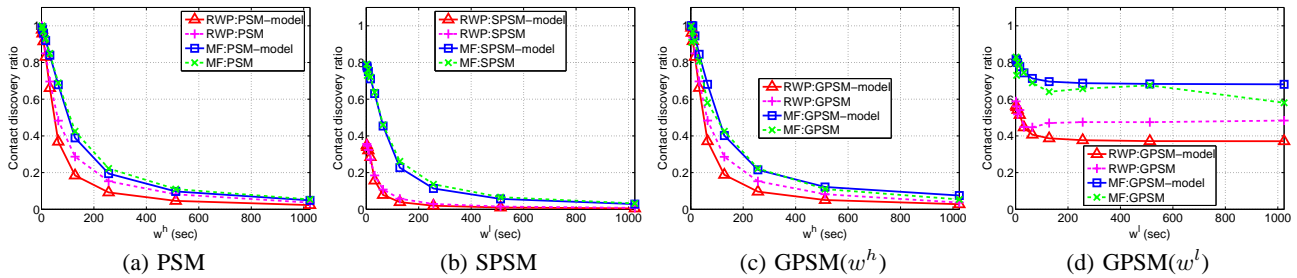


Figure 7: The actual contact discovery ratio compared to the estimated contact discovery ratio for both scenarios, MF and RWP, when wake-up intervals vary. Figure 7(c) is drawn when w^l is 1024 seconds, and Figure 7(d) is drawn when w^h is 64 seconds.

tion about traffic load helps to enhance the performance of power management mechanisms.

In Figures 8(c) and (d), we illustrate the simulation results in the MF scenario. Figure 8(c) shows that GPSM and SPSM consumes about 10% less energy than the PSM case for most of the traffic loads. In the MF scenario, ferries visit each node closely, so nodes can discover most of the contacts by the short radio range of the low-power radio while saving energy. In this case, our wake-up interval estimation for GPSM selects similar wake-up intervals for the low-power radio as SPSM and infinity for the high-power radio. As a result, SPSM and GPSM consume the least amount of energy and their delivery ratios are close to that of CAM as shown in Figure 8(d). Also, the comparison with PSM-16 shows that power management with traffic-aware optimization can save up to 80% of energy expended by PSM-16, while delivering the same number of messages.

From this evaluation, we conclude the following. First, all of the techniques save considerable energy by utilizing the traffic-aware optimization. Secondly, the performance of each mechanism depends on the node mobility model: PSM works better than SPSM when nodes pass by each other at long distances, and SPSM works better than PSM when nodes tend to meet at short distances. GPSM is flexible enough to adapt its behavior to both mobility scenarios and achieve almost the best performance in each scenario. Third, traffic-aware optimization saves energy more efficiently than pure optimization for discovery performance.

Finally, in both mobility scenarios, single-radio PSM with traffic-aware optimization achieves similar performance to dual-radio GPSM without the optimization. That is, if the additional traffic information available, the one-radio architecture could be tuned in a manner that would overcome the disadvantage of not having an extra radio.

6. RELATED WORK

There is a rich field of power management in ad-hoc networks and we point out the most relevant systems here. First, measurement studies [11, 28] show that energy consumption while idling is as high as that while receiving data. Thus, nodes can save significant energy by disabling their radios (i.e., sleeping), if not used. Exploiting this idea, many sleep/wake-up cycling mechanisms have been proposed in multihop ad hoc network to save energy while keeping network connectivity [1, 29, 33, 7, 30]. These approaches assume that a network is densely deployed, in which a node has another node within its radio range most of the time. Thus, they focus on how to overlap the time intervals in which nodes are awake to form a connected network, while allowing nodes to sleep. In our approach, we consider sparse networks in which nodes are often isolated from the rest of the network for a long time.

Secondly, significant effort has been devoted to develop energy efficient Medium Access Control (MAC) protocols [26, 31, 22, 9, 25]. These studies are initiated by observing that transmission of

multiple nodes in wireless medium interfere with each other, so at most two nodes can communicate within a radio range at a time. Thus, if a node sleeps while others are communicating, it can save energy without sacrificing its throughput. Therefore, they propose mechanisms to increase sleeping time based on traffic activity in the neighborhood. As a consequence, they are useful in a dense network rather than in a sparse network like a DTN.

Third, hierarchical power management has been implemented in various forms [24, 27, 21]. They utilize low-power devices to listen for incoming signals and to wake up the main device, if needed. Also, studies in [5, 2, 3] propose to use multiple radios to save energy by offloading data traffic in a low rate to a low-power radio or to increase the capacity of a mesh network. However, all of them assume that nodes are always close enough to other nodes. So, they did not investigate the impact of the short radio range of the low-power devices to networking performance. In this paper, we explore the limitation using simulations and propose a framework to overcome the drawbacks.

Finally, energy can be saved by using mobile nodes [23, 19, 12]. Studies in [23] utilize mobile nodes to collect data from stationary sensors. They use short range radio and low duty cycle of sensors to save energy. Also, Jun et al. [19] use a mobile node, called ferry, to route messages even in a dense network and develop power management mechanisms for mobile nodes as well as stationary nodes, in which nodes alternate between sleeping and waking up according to the predicted location of the ferry. Goldenberg et al. propose to control node mobility to optimize the network performance such as energy efficiency based on specific traffic demands while using multihop routing protocols [12]. This approach requires nodes with extra capability to control their movement. In this paper, we assume that all nodes participate in routing and do not control node mobility.

7. CONCLUSIONS

In this paper, we investigate power management in DTNs with high randomness in the node mobility. We present a hierarchical power management framework, in which nodes control two radio interfaces to discover contacts. In addition, we provide traffic-aware approximation algorithms to save energy while discovering enough contacts to deliver the traffic load in the network. Our simulation results from two mobility models show that our generalized power management mechanism could achieve better energy efficiency than mechanisms relying only on one radio for contact discovery. In addition, when the traffic load can be predicted, our approximation algorithms help nodes to save significant amount of energy while handling the expected traffic load. Finally, the additional information allows the one-radio architecture to save equivalent energy to the two-radio architecture, overcoming the disadvantage of not having the additional radio.

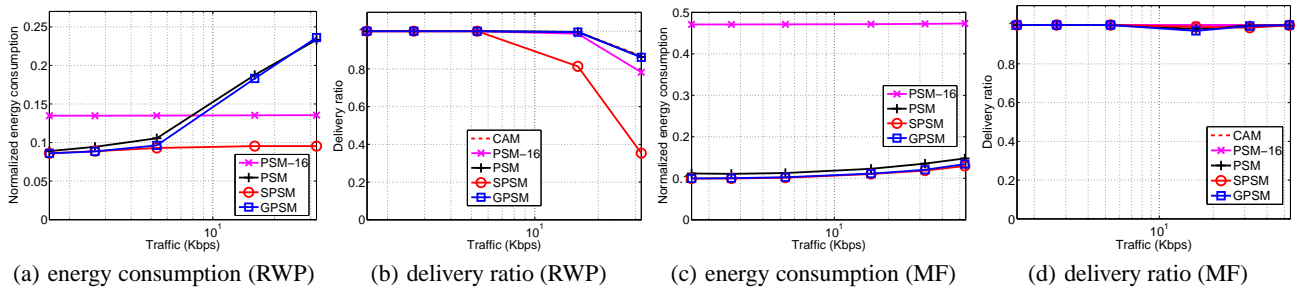


Figure 8: The impact of traffic-aware optimization to the power management in the RWP and MF scenarios

8. REFERENCES

- [1] Wireless LAN medium access control and physical layer specifications, August 1999. IEEE 802.11 Standard (IEEE Computer Society LAN MAN Standards Committee).
- [2] A. Adya, P. Bahl, J. Padhye, A. Wolman, and L. Zhou. A multi-radio unification protocol for IEEE 802.11 wireless networks. In *IEEE Broadband Networks*, 2004.
- [3] P. Bahl, A. Adya, J. Padhye, and A. Wolman. Reconsidering wireless systems with multiple radios. *ACM SIGCOMM Computer Communication Review*, 2004.
- [4] F. Bai, N. Sadagopan, B. Krishnamachari, and A. Helmy. Modeling path duration distributions in manets and their impact on reactive routing protocols. *IEEE Journal on Selected Areas in Communications*, 22(7), September 2004.
- [5] N. Banerjee, J. Sorber, M. Corner, S. Rollins, and D. Ganesan. Triage: A power-aware software architecture for tiered microservers. Technical Report 05-22, University of Massachusetts-Amherst, April, 2005.
- [6] J. Burgess, B. Gallagher, D. Jensen, and B. L. Levine. Maxprop: Routing for vehicle-based disruption-tolerant networks. In *INFOCOM*, 2006.
- [7] B. Chen, K. Jamieson, H. Balakrishnan, and R. Morris. Span: An energy-efficient coordination algorithm for topology maintenance in ad hoc wireless networks. In *MobiCom*, 2001.
- [8] Chipcon. SmartRF cc1000: Single chip very low power RF transceiver. <http://www.chipcon.com>.
- [9] W. S. Conner, J. Chhabra, M. Yarvis, and L. Krishnamurthy. Experimental evaluation of synchronization and topology control for in-building sensor network applications. In *the Second ACM International Workshop on Wireless Sensor Networks and Applications (WSNA)*, 2003.
- [10] K. Fall. A delay-tolerant network architecture for challenged internets. In *SIGCOMM*, 2003.
- [11] L. M. Feeney and M. Nilsson. Investigating the energy consumption of a wireless network interface in an ad hoc network. In *INFOCOM*, 2001.
- [12] D. K. Goldenberg, J. Lin, and A. S. Morse. Towards mobility as a network control primitive. In *Mobihoc*, 2004.
- [13] A. A. Hasson, R. Fletcher, and A. Pentland. DakNet: A road to universal broadband connectivity. Wireless Internet UN ICT Conference Case Study, 2003. <http://www.medialabasia.org/>.
- [14] S. Jain, K. Fall, and R. Patra. Routing in a delay tolerant network. In *SIGCOMM*, 2004.
- [15] D. B. Johnson and D. A. Maltz. Dynamic source routing in ad hoc wireless networks. In *Mobile Computing*, 1996.
- [16] P. Juang, H. Oki, and Y. Wang. Energy-efficient computing for wildlife tracking: Design tradeoffs and early experiences with ZebraNet. In *ASPLOS-X*, San Jose, CA, October 2002.
- [17] H. Jun, M. H. Ammar, M. Corner, and E. W. Zegura. Hierarchical power management in delay tolerant networks with traffic-aware optimization. Technical Report GT-CSS-2006-005, Georgia Institute of Technology, 2006.
- [18] H. Jun, M. H. Ammar, and E. W. Zegura. Power management in delay tolerant networks: A framework and knowledge-based mechanisms. In *IEEE Communications Society Conference on Sensor and Ad Hoc Communications and Networks (SECON)*, September 2005.
- [19] H. Jun, W. Zhao, M. H. Ammar, E. W. Zegura, and C. Lee. Trading off latency for energy in wireless ad hoc networks by message ferrying. In *IEEE PerCom International Workshop on Pervasive Wireless Networking*, March 2005.
- [20] R. H. Kravets. Cooperative solutions to the dynamic management of communication resources. In *Ph.D. thesis*, 1998.
- [21] T. Pering, V. Raghunathan, and R. Want. Exploiting radio hierarchies for power-efficient wireless device discovery and connection setup. In *VLSI Design*, 2005.
- [22] J. Polastre, J. Hill, and D. Culler. Versatile low power media access for wireless sensor networks. In *the Second ACM Conference on Embedded Networked Sensor Systems (SenSys)*, 2004.
- [23] R. C. Shah, S. Roy, S. Jain, and W. Brunette. Data MULEs: Modeling a three-tier architecture for sparse sensor networks. In *IEEE Workshop on Sensor Network Protocols and Applications (SNPA)*, 2003.
- [24] E. Shih, P. Bahl, and M. J. Sinclair. Wake on wireless: An event driven energy saving strategy for battery operated devices. In *MobiCom*, 2002.
- [25] C. Shurgers, V. Tsiatsis, S. Ganeriwal, and M. Srivastava. Topology management for sensor networks: Exploring latency and density. In *MobiHoc*, 2002.
- [26] S. Singh and C. Raghavendra. PAMAS: Power aware multi-access protocol with signalling for ad hoc networks. *ACM Computer Communication Review*, 28(3):5–26, July 1998.
- [27] J. Sorber, N. Banerjee, M. D. Corner, and S. Rollins. Turducken: Hierarchical power management for mobile devices. In *Mobisys*, 2005.
- [28] M. Stemm and R. H. Katz. Measuring and reducing energy consumption of network interfaces in hand-held devices. *IEICE Transactions on Communications*, E80-B(8):1125–1131, August 1997.
- [29] Y.-C. Tseng, C.-S. Hsu, and T.-Y. Hsieh. Power-saving protocols for IEEE 802-11-based multi-hop ad hoc networks. In *INFOCOM*, 2002.
- [30] Y. Xu, J. Heidemann, and D. Estrin. Geography-informed energy conservation for ad hoc routing. In *MobiCom*, 2001.
- [31] W. Ye, J. Heidemann, and D. Estrin. An energy-efficient MAC protocol for wireless sensor networks. In *INFOCOM*, 2002.
- [32] W. Zhao and M. H. Ammar. Message ferrying: Proactive routing in highly-partitioned wireless ad hoc networks. In *FTDCS*, 2003.
- [33] R. Zheng and R. Kravets. Asynchronous wakeup for ad hoc networks. In *MobiHoc*, 2003.