

On-demand Power Management for Ad Hoc Networks

Rong Zheng, Robin Kravets
Dept. of Department of Compute Science
University of Illinois at Urbana-Champaign
Urbana, IL 61801
E-mail: {zheng4,rhk}@cs.uiuc.edu

Abstract—Battery power is an important resource in ad hoc networks. It has been observed that in ad hoc networks, energy consumption does not reflect the communication activities in the network. Many existing energy conservation protocols based on electing a routing backbone for global connectivity are oblivious to traffic characteristics. In this paper, we propose an extensible on-demand power management framework for ad hoc networks that adapts to traffic load. Nodes maintain soft-state timers that determine power management transitions. By monitoring routing control messages and data transmission, these timers are set and refreshed on-demand. Nodes that are not involved in data delivery may go to sleep as supported by the MAC protocol. This soft state is aggregated across multiple flows and its maintenance requires no additional out-of-band messages. We implement a prototype of our framework in the ns-2 simulator that uses the IEEE 802.11 MAC protocol. Simulation studies using our scheme with the Dynamic Source Routing protocol show a reduction in energy consumption near 50% when compared to a network without power management under both long-lived CBR traffic and on-off traffic loads, with comparable throughput and latency. Preliminary results also show that it outperforms existing routing backbone election approaches.

I. INTRODUCTION

With the proliferation of portable computing platforms and small wireless devices, ad hoc wireless networks have received more and more attention as a means for providing data communications among devices regardless of their physical locations. Wireless communication has the advantage of allowing untethered communication, which implies reliance on portable power sources such as batteries. However, due to the slow advancement in battery technology, battery power continues to be a constrained resource and so power management in wireless networks remains to be an important issue.

It has been observed that in ad hoc networks, energy consumption does not always reflect active communication in the network [1]. Experimental results reveal that the energy consumption of wireless devices in an idle state is only slightly smaller than that in a transmitting or receiving state. Therefore, it is in general desirable to turn the radio off when it is not in use. Motivated by these observations, several energy conservation protocols [2], [3] have been proposed to take advantage of route redundancy in dense ad hoc networks by turning off devices that are not required for global network connectivity. However, in these protocols, the decision about which set of

nodes to leave on is only based on geographical/topological information, thus is oblivious to the actual traffic load in the network. Since many applications of ad hoc networks are data-centric, maintenance of global connectivity is costly and unnecessary when no traffic or only localized traffic is present in the network.

Various techniques, both in hardware and software, have been proposed to reduce energy consumption for mobile computing devices in wireless LANs [4], [5]. In contrast, power management in ad hoc networks is a more difficult problem for two reasons. First, in ad hoc networks, a node can be both a data source/sink and a router that forwards data for other nodes and participates in high-level routing and control protocols. Additionally, the roles of a particular node may change over time. Second, there is no centralized entity such as an access point to control and maintain the power management mode of each node in the network, buffer data and wake up sleeping nodes. Therefore, power management in ad hoc networks must be done in a distributed and cooperative fashion. A major challenge to the design of a power management framework for ad hoc networks is that energy conservation usually comes at the cost of degraded performance such as lower throughput or longer delay. A naive solution that only considers power savings at individual nodes may turn out to be detrimental to the operation of the whole network.

In this paper, we propose an on-demand power management framework targeting generic ad hoc networks. To achieve reduced energy consumption while maintaining effective communication, our framework integrates routing information from on-demand ad hoc routing protocols and power management capabilities from the MAC layer. Energy conservation is achieved by judiciously turning on and off the radios of specific nodes in the network. The novelty of our framework is that such power management decisions are driven by active communications in the network. For the purpose of energy conservation, connectivity is only maintained between pairs of senders and receivers and along the route of data communication.

Transitions between power management modes for each node are associated with a soft-state timer that is established and refreshed by data and control messages in the network. Once the soft state is established, subsequent data delivery can

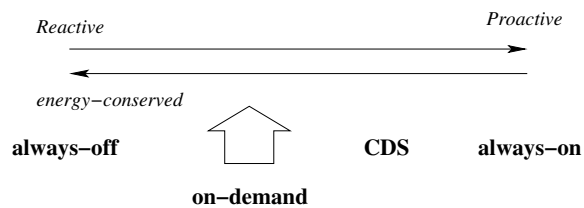


Fig. 1. Design space of power management schemes

be expedited without incurring additional delays from waking up sleeping nodes along the route. The length of the soft-state timer reflects the adaptiveness of the power management framework to variations in traffic load. Since the operations of transmitting to a sleeping node and an active node are different, we present mechanisms to discover a neighbor's power management mode. In this context, neighbor discovery is challenging because a node in power-save mode cannot monitor the channel consistently. Therefore, any neighbor information may be ambiguous. This situation is even worse if nodes are mobile.

Our framework is not limited to any specific routing or MAC protocols. This extensibility is a key benefit of our design since it enables the use of our framework in various scenarios and allows the integration of new protocols as they become available. To verify our framework, we present a prototype using the IEEE 802.11 MAC protocol and evaluate it using Dynamic Source Routing (DSR) [6] and greedy geographical forwarding protocol in the ns-2 [7] simulator. Under a wide range of traffic patterns and load, our prototype achieves 40-60% savings in power consumption as compared to a network without power management. In addition, our prototype minimally increases latency during the initial setup stage, but achieves an average latency comparable to a network without power management.

The rest of the paper is organized as follows. We first layout the design space for power management protocols in ad hoc wireless networks and give a brief overview of existing approaches in Section II. Then we discuss how each approach fits into the design space. In Section III, we present the building blocks and technical details of our on-demand power management framework. Section IV describes a prototype based on IEEE 802.11 MAC. Extensive simulation results are presented in Section V. Finally, we conclude the paper and discuss future extensions in Section VI.

II. DESIGN SPACE

Power management in ad hoc networks spans all layers of the communication protocol stack. Each layer has access to different types of information about the communication in the network, and thus uses different mechanisms for power management. The MAC layer does power management using local information while the network layer can take a more global approach based on topology or traffic characteristics. In this paper, we consider power management approaches that save energy by turning off the radios of nodes in the network. Other

energy conservation mechanisms such as topology control and power controlled MAC protocols [8], [9], [10] are considered orthogonal and the benefits can be combined.

Similar to ad hoc routing protocols, power management schemes range from proactive to reactive. The extreme of proactive can be defined as *always-on* (i.e. all nodes are active all the time) and the extreme of reactive can be defined as *always-off* (i.e. all nodes are in power saving mode by default) (see Figure 1). Given the dynamic nature of ad hoc networks, there needs to be a balance between proactiveness, which generally provides more efficient communication, and reactiveness, which generally provides better power saving.

In this section, we outline the design space of power management in ad hoc networks and describe where existing approaches fit into this design space based on their adaptability to network traffic.

MAC Layer approaches: At the MAC layer, power management decisions are made based on local information. The time scale for power management can be per-packet or a short time interval. Such approaches are limited by the lack of access to information about the topology and traffic in the network.

The PAMAS power-saving medium access protocol [11] turns off a node's radio when it overhears a packet not addressed to it. The effectiveness of PAMAS is limited to reducing the power consumption of processing unnecessary packets. Note that PAMAS alone can be considered a proactive approach to power management, however it may be combined with most high level power management schemes that aim to reduce idle time energy consumption.

The IEEE 802.11 MAC provides low-level support for power management such as buffering data for sleeping nodes and synchronizing nodes to wake up for data delivery. The network interface has five physical states: transmitting, receiving, idle, sleeping and completely power-off. Energy consumption in the sleeping state is significantly less than in the transmitting/receiving/idle state. In the IEEE 802.11 specification, a node can be in one of two power management modes, active mode (AM) or power-save mode (PS). In active mode, a node is awake and may receive data at any time. In power-save mode, a node wakes up periodically to check for incoming traffic. The transition between power management modes is left to higher-level power management protocols and is unspecified in the documentation.

STEM [12] proposes a similar approach to the IEEE 802.11 power management, but uses an independent control channel to avoid the clock synchronization needed by IEEE 802.11. STEM uses asynchronous beacon packets in a second control channel to wakeup intended receivers. After transmissions have ended (e.g. after a timeout, etc.), the node turns its radio off in the data channel. Similar to IEEE 802.11, sleeping nodes with traffic destined for them are woken up on demand, but decisions about when a node should go back to sleep are based on local information. STEM does not provide mechanisms for indicating the power management state of a node. Instead, the power management state is only maintained on a per-link

basis between nodes with active communication. Therefore, it is possible that an initiator node will experience the delay from waking up a receiver node, even if the receiver is already awake due to recent communication with a third node.

In S-MAC [13], the authors propose a mechanism called message passing that modifies a network allocation vector (NAV) for virtual channel reservation in IEEE 802.11 MAC type of protocols. The length of NAV is determined by the duration of a burst of messages. The virtual reservation serves two purposes: (1) it mandates the receiver to remain on throughout the transmission of the burst, and (2) it prevents other nodes from transmitting during this interval. Though message passing may be desirable for certain types of applications for sensor networks, it can be inefficient for more generic ad hoc networks. Additionally, the reservations may cause some nodes to be starved.

A pure MAC layer approach as specified by the IEEE 802.11 MAC (i.e. nodes are always in power-save mode) can be considered as the most reactive approach to power management in our design space. In Section V, we demonstrate that a network that relies solely on the IEEE 802.11 MAC for power management can be highly inefficient even though some communication is still possible. As future research, we will investigate the interaction between intelligent MAC layer approaches, such as STEM and S-MAC, with our on-demand framework.

Connected Dominating Set Approaches: At the network layer, power management schemes can take advantage of topological information. The connected dominating set approaches use neighborhood or global information to decide the set of nodes that form a connected dominating set (CDS) for the network, where all nodes are either a member of the CDS or a direct neighbor of one of the members. Nodes in the CDS serve as the "routing backbone" and remain on all the time to maintain global connectivity. All other nodes can choose to sleep if necessary.

CDS approaches such as GAF [3] and SPAN [2] conserve energy by reducing routing redundancy in dense networks. Selection and maintenance of the CDS requires local broadcast messages that may consume a significant amount of energy [1]. In addition, regardless of whether or not traffic is present in the network, all backbone nodes must be on all the time. Therefore, CDS approaches can be categorized as proactive.

Based on these observations about MAC layer and CDS approaches, we propose an on-demand power management framework to explore the design space between proactive and reactive by adapting to the traffic characteristics inside the network.

III. ON-DEMAND POWER MANAGEMENT

The goal of on-demand power management is to base power management decisions on traffic patterns in the network. By reacting to changes in these patterns, nodes that do not carry any traffic can be dispensed from consuming a significant amount of energy. Varying the adaptiveness to network load in

our protocol balances the trade-off between latency, throughput and energy consumption.

The key idea of our on-demand power management framework is that transitions from power-save mode to active mode are triggered by communication events such as routing control messages or data packets and transitions from active mode to power save mode are determined by a soft-state timer. The soft-state timer is refreshed by the same communication events that trigger a transition to active mode. A node keeps track of its neighbors' power management mode either by HELLO messages or by snooping transmissions over the air. For direct unicast messages, if the next hop is in active mode, the message is delivered immediately as allowed by the queuing discipline.

A. A cross-layer design for power management

Power management in ad hoc networks can benefit from a cross-layer design that leverage both network layer and MAC layer information. Knowledge about route setup and packet forwarding can provide hints about when power management should be performed. Since the route discovery phase of on-demand routing protocols determines the path subsequent packets will follow, nodes along this route should be as responsive as possible. On the other hand, any effective power management protocol requires a mechanism to awaken a sleeping receiver when packet delivery is imminent. This is usually handled by low-level mechanisms at the MAC or physical layers. Higher-level power management techniques can benefit from information about and access to the mechanisms used to provide such services.

Our power management framework leverages the capability of modern MAC protocols, such as the IEEE 802.11 MAC, to switch power management states of nodes and buffer data if necessary for sleeping nodes. It also uses routing information to decide when to turn nodes on and off, which ties energy consumption with active communication in the network.

B. Power management mode and state transition

In our framework, a node can be in one of two power management modes: active mode (AM) and power-save mode (PS). In active mode, a node is awake and may receive data at any time. In power-save mode, a node is sleeping most of the time and wakes up periodically to check for pending messages. Packets destined to a node in power-save mode will experience delay on the order of the length of the sleeping cycle.

Transitions from power-save mode to active mode are triggered by communication events in the network. Transitions from active mode to power-save mode are determined by a soft-state keep-alive timer. Initially, all nodes are in power-save mode. Upon reception of packets, a node starts the keep-alive timer and switches to active mode. Timer values depend on the type of packet received. Upon expiration of the keep-alive timer, a node switches from active mode to power-save mode.

If all packets trigger a node to stay awake with a keep-alive timer on the order of the network lifetime, our scheme degenerates to an always-on network without power management.

On the other hand, if the keep-alive timer is always set to zero, our framework degenerates to the most reactive MAC layer approach discussed in Section II. Therefore, the choice of different keep-alive timer values varies the reactivity of the protocol and strikes different trade-offs between energy consumption and data delivery efficiency.

In an ad hoc network, if a path is going to be used, the nodes along that path should be awake as to not cause unnecessary delay for data transmission. If a path is not going to be used, the nodes should be allowed to sleep. During the lifetime of the network, different messages will indicate different levels of "commitment" to using a path. Knowledge of the semantics of such messages can help make better power management decisions, which is a missing piece in most MAC layer power management approaches.

On one end, most control messages (e.g. link state in table-driven ad hoc routing protocols, location updates in geographical routing, route request messages in on-demand routing protocols etc.) are flooded throughout the network and provide poor hints for the routing of data transmissions. Such control messages should not trigger a node to stay in active mode. On the other end, data transmissions are usually bound to a path on relatively large time scales. Therefore, data transmissions are a good hint for guiding power management decisions. For data packets, the keep-alive timer should be set on the order of the packet inter-arrival time to ensure that nodes along the path do not go to sleep during active communication. There are also some control messages, such as route reply messages in on-demand routing protocols and query messages in sensor networks, that provide a strong indication that subsequent packets will follow this route. Therefore, such messages should trigger a node to switch to active mode. The time scale of the keep-alive timer for such a transition should be on the scale of the end-to-end delay from source to destination so the node does not transition back to power save mode before the first data packet arrives.

One important feature of the keep-alive timer is that it is refreshed on demand. Whenever a node receives a routing message or a data packet, it sets the timer with the maximum of what is left for the current keep-alive timer and the value associated with the received message. Therefore, only per-node, instead of per-flow, information is needed for power management. Ideally, the keep-alive values for data packets should be larger than the inter-arrival time of data packets. In reality, since a node can be both data source/sink and forwarding router simultaneously, its keep-alive timer is an aggregation of various timer values, i.e., old timer will be extended when new communication events arrive. Therefore, the performance is quite insensitive to the choice of these timer values.

C. Obtaining neighbors' power management mode

Since communication with a neighbor is only possible if the neighbor is in active mode, it is necessary for nodes to track power management modes of neighbors. In our framework, each node maintains a neighbor list that caches a neighbor's

power management mode and a time-stamp of the most recent update from this neighbor.

A neighbor's power management mode can be discovered in two ways. The first way is through explicit local HELLO message exchanges with piggybacked information about the power management mode of a node. HELLO messages should be transmitted at fixed intervals regardless of the power management mode of a node. Link failure is assumed if no HELLO messages have been received during successive intervals, since the loss of only one HELLO message may have been caused by a broadcast collision.

The second way to discover a neighbor's power management mode is via passive inference. Depending on the capability of the hardware and the MAC protocol, a node may be able to operate in promiscuous mode and passively snoop messages in the air. With MAC layer support, a node's power management mode can be piggybacked in the control header of MAC layer data units. There are two challenges to using passive inference. First, nodes in power-save mode cannot hear messages from their neighbors and so do not have a good basis for determining the power management mode of their neighbors. Second, nodes in power-save mode may not be transmitting and so their neighbors will have difficulty differentiating nodes that are in power-save mode from nodes that are away or dead. Therefore, special care must be taken to distinguish between nodes that move away from ones that are in power-save mode.

Since the use of HELLO messages is expensive, our framework uses two types of indicators for such passive inference. The first indicator is a lack of communication during a time interval. When no communications have been observed from a node that was in active mode, the neighbor is assumed to be in power-save mode. The value of this interval should be based on the keep-alive timer since the length of keep-alive timer indicates the maximum amount of time a node commits to be in active mode when no messages are received. If a node does not hear from its neighbor during the keep-alive period, it is very likely that either that neighbor moved away or it has switched to power-save mode.

The other indicator is packet delivery failure to the neighbor (e.g. indicated by an RTS retry time out in IEEE 802.11). Based on the observed power management mode of the neighbor, a packet delivery failure is treated in two stages. First, if the neighbor was originally in active mode, it is considered to have switched to power-save mode. Second, if the neighbor was originally in power-save mode, it is now considered unreachable. Data for this node is discarded at the MAC layer. The rationale for this two-stage process is that transition to an intermediate stage provides a second chance to salvage data for a neighbor that has switched to power-save mode since the last update.

Compared to using HELLO messages, passive inference does not rely on additional control messages, which is more desirable from an energy conservation perspective. However, the ambiguity of link failure and the power management mode of a neighbor can result in delayed data transmission, but as

will be shown in the Section V, this approach in general works well.

IV. A PROTOTYPE BASED ON THE IEEE 802.11 MAC

In this section, we present a prototype of our framework based on the IEEE 802.11 MAC. First, we give a brief overview of IEEE 802.11 power management functions and then we discuss the implementation details of our prototype. Note that our on-demand framework can be easily implemented over other MAC protocols including those using asynchronous wakeup mechanisms [14].

A. Overview of IEEE 802.11 Power Management in ad hoc networks

In the IEEE 802.11 specification, all nodes in the network are synchronized to wake up periodically. Broadcast/multicast messages or unicast messages to a power-saving node are first announced during the period when all nodes are awake. The announcement is done via an ad hoc traffic indication message (ATIM) inside a small interval at the beginning of the beacon interval called the ATIM window.

During the ATIM window, nodes that have buffered data for sleeping nodes transmit an ATIM management frame that contains the identity of the intended receivers. If a node receives a directed ATIM frame during the ATIM window (i.e. it is the designated receiver), it sends an acknowledgment and stays awake for the entire beacon interval waiting for the data to be transmitted. Broadcast/multicast messages announced in the ATIM need not be acknowledged. Immediately after the ATIM window, nodes can transmit buffered broadcast/multicast frames, data packets and management frames addressed to nodes that have acknowledged a previously transmitted ATIM frame. Following the transmission of all buffered data, nodes transmit data destined to nodes that are known to be in the active state for the current beacon interval. In IEEE 802.11, a node's power management status is indicated in the frame control field of the MAC header for each packet.

B. Our prototype

We implemented a prototype of our framework based on the IEEE 802.11 MAC. For routing, we assume an on-demand routing protocol such as DSR [6] or AODV [15]. We also experiment with stateless routing protocols such as GPSR and achieve similar results. Due to space limitations, these results have been omitted but can be found in [16].

The complete state transition diagram is shown in Figure 2. Transitions between power-save and active mode are triggered by packet arrivals and expiration of the keep-alive timer. Sub-state transitions inside power-save or active mode indicate the physical state of the node and are controlled by the IEEE 802.11 MAC power management functions.

To maintain the neighbor list, the prototype uses passive inference to update neighbors' power management modes and link states. Nodes snoop transmissions in the air when they are awake and update their neighbor lists based on the control field of the MAC header of packets. Entries for

unreachable neighbors are purged periodically. Due to the use of beacon messages, changes in link availability can be detected proactively as follows.

Let the beacon interval be I and let the degree of a node be bounded by d . The interval N (measured in units of beacon intervals) of two successive beacon messages sent by node n follows the geometric distribution $N \sim \frac{1}{d}(1 - \frac{1}{d})^{N-1}$ with mean d . Therefore, if a node has not heard any beacon messages from a particular neighbor for more than $c \cdot d$ beacon intervals, where c is a protocol-specific constant, the node is likely away or "dead". The degree of a node is obtained from the neighbor list and a node can use its own degree to approximate its neighbor's degree. This method can be combined with events of packet delivery failure to better infer the availability of a link between neighbors. The major benefit of inferring a neighbor's state by snooping is that it does not incur out-of-band control messages and therefore scales to large networks.

To determine the values of various keep-alive timers, consider a k -hop route from source n_0 to sink n_k . Suppose the one-way delay from node n_0 to node n_i on this route is d_i . The inter-arrival time of data packets at the source is $1/\lambda$. Let the beacon interval be I . Therefore, the time it takes for the route request message to reach node n_i from n_0 is $i * I + d_i$ on average. The time it takes node n_i to receive the route reply message is $(2k + 1 - i) * I + d_{k+1} + d_i$ under the assumption of symmetric routes. Assuming the data source will immediately transmit the data upon reception of the route reply message, the time between the reception of the route reply message and the reception of the first data packet at node n_i is $i * I + 2 * d_i$. Finally, assuming no additional queuing at intermediate nodes, the packet inter-arrival time at node i is $1/\lambda$. Therefore, on the pessimistic side, the length of the keep-alive timers for different message types should be chosen to be larger than these estimates to ensure low latency for data delivery. If available, information about the network dimensions and traffic patterns can be used to select these values based the above discussion. In our performance evaluation in Section V, we do not assume availability of such knowledge. In our implementation, we set RTRQ_KEEPALIVE to 0, RTRL_KEEPALIVE to 5 seconds, DATA_KEEPALIVE, SRC_KEEPALIVE and DST_KEEPALIVE to 2 seconds. The REFRESH_INTERVAL is set to 5 seconds. Since keep-alive timer values can be aggregated, the performance is quite insensitive to the choice of these values. As part of our future work, we will investigate techniques to adapt keep-alive timer values based on measurements in the network.

V. PERFORMANCE EVALUATION

We implemented our prototype in the ns-2 [7] network simulator using the CMU wireless extension [17]. To evaluate the effectiveness of our proposed scheme, we conducted several simulations using different traffic models in both static and mobile networks.

The effectiveness of power management schemes can be evaluated by (1) longevity: the network should remain op-

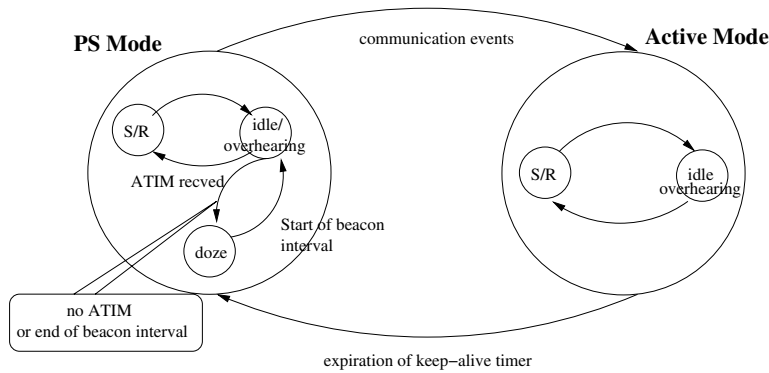


Fig. 2. State transition diagram for power management modes enhanced with IEEE 802.11 physical states

erational for as long as possible, and (2) efficiency: data transmissions should experience low loss and low latency. Longevity is normally characterized by the lifetime of the network, which is application-specific and is tightly coupled with how the network is being used. In this paper, we only focus on the power consumption per unit data delivery defined as follows.

$$energy_goodput = \frac{total_bit_transmitted}{total_energy_consumed} \quad (1)$$

where the total bits transmitted are calculated for *application-layer data packets* only. The unit of energy goodput is bit/J, which in essence captures the energy utilization of the network with all control overhead considered. Efficiency of data delivery is characterized by the end-to-end latency and the packet delivery ratio defined as the total amount of data received divided by the total amount of data transmitted.

For comparison, we use the most reactive and proactive schemes as baselines. The first is the pure IEEE 802.11 MAC layer power management, termed as *always-off* (i.e. every node is always in power-save mode). The second is without any power management, termed as *always-on* (i.e. all nodes are active throughout the simulation). Unless otherwise stated, there are 50 nodes randomly placed in a 1500mx300m rectangular plane. All nodes communicate with half-duplex wireless radios that conform to IEEE 802.11-based WaveLAN wireless radios with a bandwidth of 2Mbps and a nominal transmission radius of 250m. In all simulation scenarios, the network is never partitioned and there are no error-induced losses. DSR is used for routing. Similar experiments were performed with AODV with similar results. We use the same energy model as in [2], which is shown in Table I. The energy consumption for switching between awake and sleeping states is negligible and thus not considered here. All data packets are of length 128 bytes. Different data packet sizes will affect the throughput of all schemes in a similar fashion. The beacon interval and ATIM window are set to 0.4s and 0.02s respectively.

TABLE I
POWER CONSUMPTION MODEL

| Transmit | Receive | Idle | Sleep |
|----------|---------|-------|-------|
| 1400mW | 1000 mW | 830mW | 130mW |

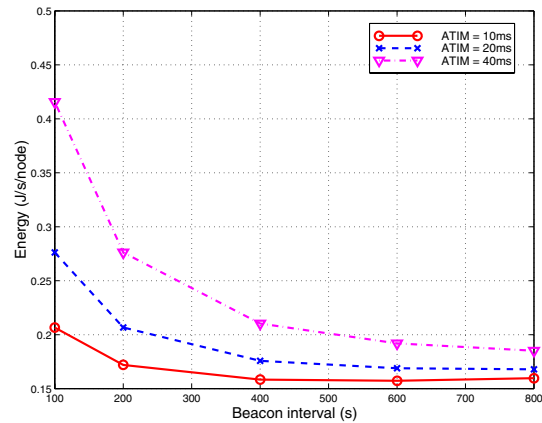


Fig. 3. Energy consumption of beacon messages

A. Effectiveness of the IEEE 802.11 power management functions

The transmission of beacon messages consumes power. Furthermore, regardless whether there is an announcement or acknowledge frame, a node must be awake during the ATIM window. Therefore, it can be expected that the smaller the beacon interval, the more power will be consumed. Similar arguments apply to the length of the ATIM window.

Figure 3 plots the energy consumption of beacon messages against the beacon interval and ATIM window size for a static 1000mx1000m network with 75 nodes. This measurement indicates how much “raw” energy is consumed in a network implementing IEEE 802.11 power management without any packet delivery. From Figure 3, we can roughly tell that once the ratio of the ATIM window size and the beacon interval is fixed, the power consumption is almost constant. This is because in the power model we use, the difference between the power consumption of transmitting, receiving and idle states

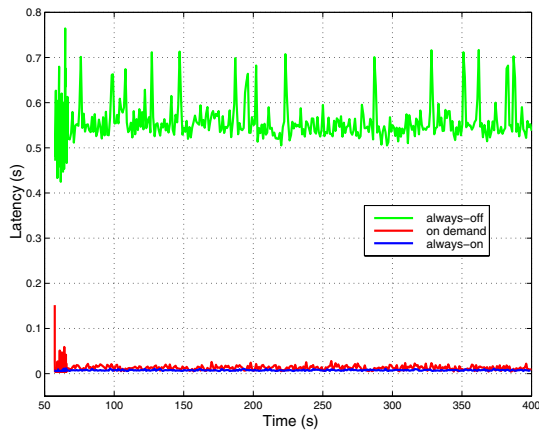


Fig. 5. End-to-end delay of one 3-hop connection, rate = 1kbps

is not very significant. This implies that to reduce the beacon interval while keeping a roughly constant power consumption for beacon messages, the length of ATIM window should be reduced by the same ratio.

B. Long-lived CBR traffic in static networks

Next, we study the performance of our prototype in a static network. In this set of simulations, we simulate long-lived CBR connections at different transmission rates. We compare loss rate, latency and energy consumption for an *always-on* scheme, an *always-off* scheme and our on-demand prototype. There are 10 randomly chosen sender-receiver pairs started randomly between 0 and 100s. The simulation results presented here are averages over four different scenarios.

Figure 4 shows the packet delivery ratio and energy goodput as traffic load changes. The *always-off* scheme does not work well under high traffic load. The packet delivery ratio decreases drastically as the load increases. When the traffic load is high, it is possible that there is not enough time to announce all buffered data packets (and get acknowledgments back) due to the limit of the ATIM window size. Depending on the implementation of the MAC and routing protocols, packets get dropped due to delayed transmission and incur further re-transmissions or route discovery. This “chain effect” will result in a pathological network.

As shown in Figure 4, our prototype achieves similar packet delivery ratios to the *always-on* scheme, while the energy throughput is nearly doubled. The reduced packet delivery ratio at higher traffic load is due to the fact that data delivery is only possible outside the ATIM window since data arriving during the ATIM window need to be buffered temporarily. This slightly lowers the capacity of the network. At low traffic load, the *always-off* scheme can achieve higher energy goodput than the *always-on* scheme, but as the traffic load increases, the energy goodput will eventually suffer due to high loss. In comparison, our on-demand prototype works consistently better than both schemes. The linear area of the energy consumption curve corresponds to the region in which the energy consumption is constant (independent of traffic

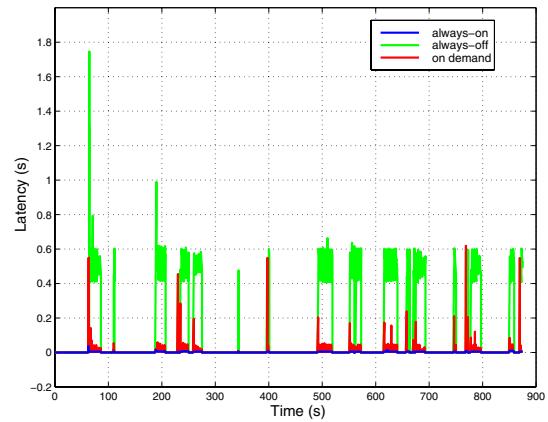


Fig. 7. End-to-end delay of one connection under on-off traffic, rate = 1kbps

load). When the traffic load is low to medium, no matter how fast sources transmit, the percentage of time a node stays awake in both our prototype and the *always-on* scheme is roughly constant since there is always data to deliver. Our prototype gains by reducing the number of waking nodes. As the traffic load increases, the number of collision gets higher and the growth of energy goodput slows down for all schemes.

Figure 5 shows the end-to-end packet delay for a single 3-hop connection during the simulation. With our scheme, apart from the initial setup stage, data packets experience similar latency as those in the *always-on* scheme. This is because after the initial setup stage, nodes in active mode have been established along the route. A node can deliver unicast data directly to its next hop neighbor without the need to make an announcement in the ATIM window.

C. On-off traffic

To understand how well our proposed protocol works under more realistic traffic patterns, we simulate on-off traffic with 30 sender-receiver pairs. Both busy and idle intervals follow exponential distribution with means of 10s and 100s respectively. The simulations run for 900s. Note that with 30 sender-receiver pairs, most of the 50 nodes in the network are either involved in data forwarding, sending or reception at some time in the simulation.

Similar to the previous set of simulations, we compare the packet delivery ratio and energy goodput vs. traffic load for different schemes. Again, we observe high packet delivery ratio (or low loss rate) and high energy goodput for our prototype (see Figure 6). The end-to-end packet delay of a 3-hop connection is shown in Figure 7. Since the idle period of the connection is very long, the keep-alive timers at intermediate nodes will eventually time out. Nodes must be woken up at the beginning of the next busy period. This explains the spikes at the beginning of each busy interval.

One thing noticeable about the energy goodput curves in Figure 6 is that under very light traffic load, the *always-off* scheme achieves slightly higher energy goodput than our prototype. This is because, in the IEEE 802.11 MAC, a node

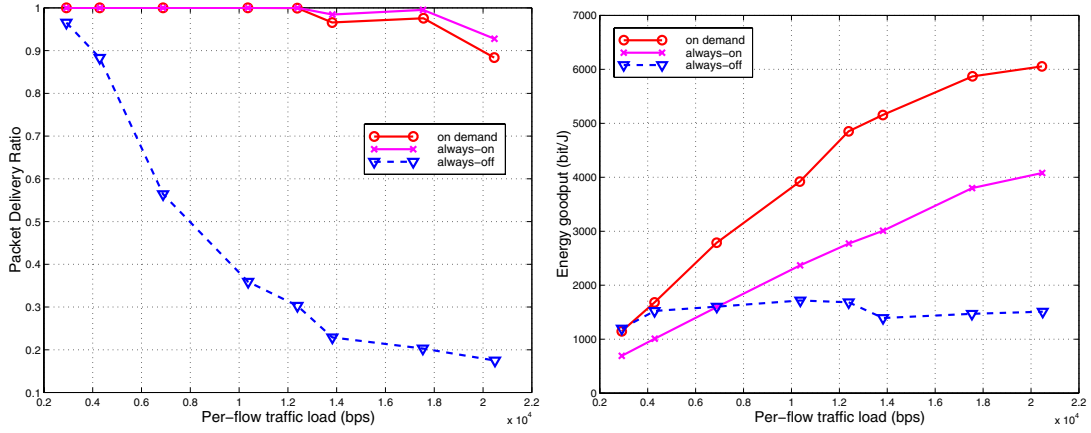


Fig. 4. Packet delivery ratio and energy goodput vs. pause time, 10 long-lived CBR connections, 50 nodes, 1500x300 static network

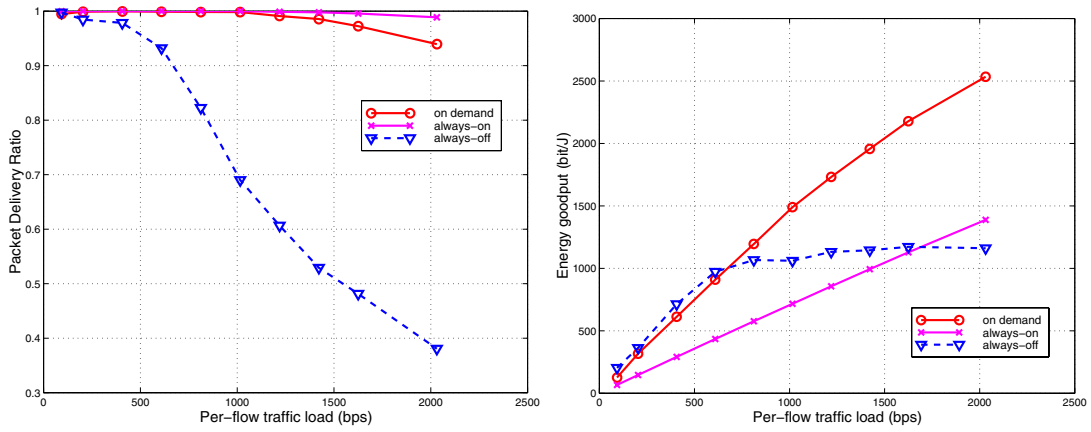


Fig. 6. packet delivery ratio and energy goodput vs. traffic load. 30 on-off connections, on = 10s, off = 100s, 50 nodes, 1500x300 static network

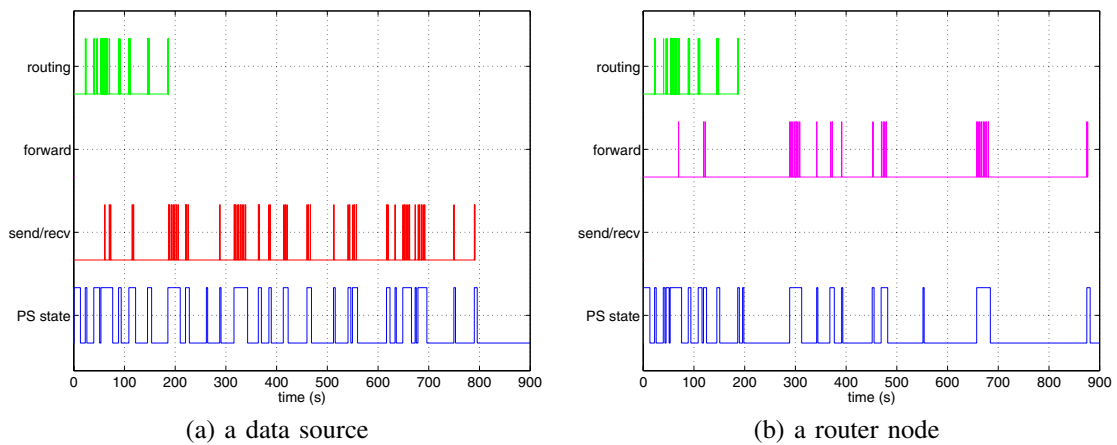


Fig. 8. Actual duty cycles of a node

will go to sleep at the end of ATIM window if no ATIM announcement is received. At light traffic loads, for example, when the sending rate in busy periods is 1kbps, the inter-arrival time of data packets is roughly 1 second, which spans several beacon intervals. The *always-off* scheme allows further energy savings by turning off nodes on a time scale as small as a beacon interval. Our prototype only switches power management modes on the time scale of the keep-alive timers. To better illustrate this point, we compare the time for actual packet delivery to the time that the wireless interface is turned on in the simulation. For clarity of presentation, we filter out the jumps due to beacon messages since they are very short and frequent.

Figure 8 shows the duty cycle of a data source/sink node and a forwarding node. We separate routing messages from data packets. A ‘1’ in the top three plots corresponds to an arrival or departure event of the corresponding type of packet. The bottom plot shows the time intervals of the node’s active mode with a ‘1’ corresponding to active. This plot consists of an envelope of the union of the top three plots. The values of DATA_KEEPAALIVE and RTRL_KEEPAALIVE determine the tail of the envelop. Since the start time of each sender-receiver pair varies, there are routing discovery/reply messages up to 200 seconds into the simulation. Since there is no mobility in this network, all packet losses are caused by collisions. Recall in Section III, we described the two-stage process to determine link availability to a neighbor. In the simulation, no additional route discovery is incurred after 200s. This indicates that the two-stage process works well in distinguishing a unreachable node from a neighbor in power-save mode.

D. Impact of mobility

In this section, we study how mobility affects the performance of our on-demand prototype. With mobility, nodes on inactive routes may still remain on for the rest of their keep-alive timer. As a result, we would expect that the gain in energy goodput will be reduced in the case of high mobility.

The simulation setup is as follows. The maximum speed of each node is 20m/s and the pause time varies from 15s to 75s. The results presented are an average of four different scenarios of 400-second runs.

As shown in Figure 9, the energy savings of our proposed scheme is not as significant as in the static scenarios. However, it still conserves a significant amount of energy compared with the *always-on* scheme. Also it performs much better than the *always-off* scheme in terms of both energy goodput and packet delivery ratio. An *always-off* network is no longer functional in high mobility with loss rates as high as 50-60%. An interesting observation is the *always-off* scheme does not conserve any energy in the case of high mobility. The reason is that frequent route discovery messages flooded in the network cause a node to stay awake most of the time.

Figure 10 shows the end-to-end delay of a 3-hop connection. When a link fails due to mobility, route discovery messages will be sent out. Some nodes on the new routes might be in

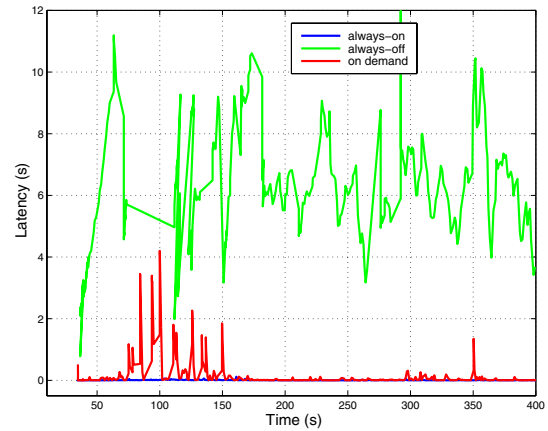


Fig. 10. End-to-end delay for one connection, speed = 20m/s, pause = 50s, rate = 1kbps

power-save mode and need to be triggered into active mode, causing the spikes in the on-demand curve.

E. Comparison with GAF

It is very difficult if not impossible to make a fair comparison between our framework and existing CDS approaches like Span and GAF, since the different schemes are based on different assumptions. We highlight the major difference between our framework, GAF and Span in Table II.

Both Span and GAF assume that data sinks and sources are separated from pure forwarding nodes in their evaluation. In the case of mixed source/sink/forwarding nodes scenarios as used in the previous simulations, the specification of both protocols is incomplete. GAF has no mechanism for signaling the data sink for incoming data. In Span, it is unclear whether the election of coordinators should consider the fact that some nodes may be required to be turned on as data sources or sinks.

In this section, we compare our prototype with GAF since it can be readily used with any routing protocol. To avoid the need to signal data sinks for incoming packets in GAF, we only simulate scenarios where data sources/sinks are at the periphery of the network and internal nodes are dedicated for routing. We simulate a 1000mx1000m plane with 85 nodes. Connections are between nodes n_{2i} and n_{2i+1} where $i = 0, 1, \dots, 4$. The locations of nodes n_{2i} and n_{2i+1} are $(0, i*250)$ and $(1200, i*250)$ respectively. Data sources/sinks do not participate in packet forwarding. In GAF, the data sources/sinks are on all the time. We simulate on-off CBR traffic with a fixed busy interval of 10s and an idle interval varying from 0s to 200s. During a busy interval, each source is transmitting at 1kbps. We use the same power consumption model for both schemes as listed in Table I. There is no mobility in this network. For comparison, we also simulate the *always-on* scheme.

Figure 11 shows the energy consumption of routing nodes under the different schemes. The energy throughput metric does not apply to this set of simulations since the energy consumption of senders and receivers are not evaluated in

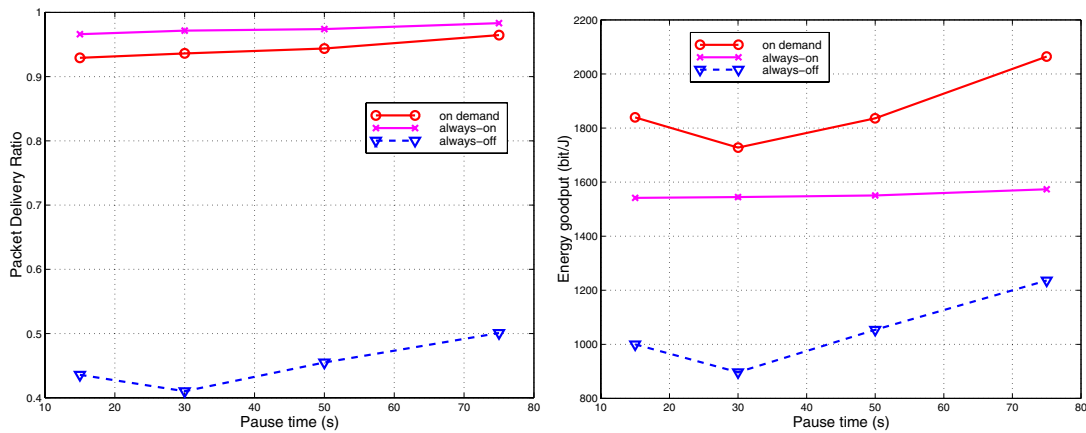


Fig. 9. Packet delivery ratio and energy goodput vs. traffic load with mobility, 10 CBR conn., 50 nodes, 1500mx300m region, speed = 20ms

TABLE II
QUALITATIVE COMPARISON OF VARIOUS POWER MANAGEMENT PROTOCOLS

| | MAC Support | GPS | Routing | Separation of data sink/source and routers | Local broadcast | Reactive |
|-----------|-------------|-----|---------|--|-----------------|----------|
| GAF | no* | yes | any | yes | yes | no |
| Span | yes | no | any** | yes | yes | no |
| On demand | yes | no | any | no | no | yes |

* Conceptually, if GAF is applied to networks with mixed data sink/source and routing nodes, signaling mechanisms are needed to wake-up data sinks upon the arrival of data packets. One possible solution is the use of the IEEE 802.11 MAC.

** The current implementation of Span in ns-2 is coupled with geographical routing protocols.

GAF. The energy consumption of both GAF and the *always-on* scheme are roughly constant with respect to the idle interval. Regardless of whether there are data transmissions in the network, either the routing backbone or all the nodes need to be on all the time. By comparison, our prototype achieves further energy conservation by adjusting a node's duty cycle based on its traffic load. The average packet delivery ratio for GAF, on-demand and the *always-on* scheme are 84%, 99.4% and 100% respectively. The reason that the packet loss rate is higher in GAF is that the rotation of grid leaders will induce packet losses.

One interesting data point in Figure 15 is energy consumption when the idle period is 0. This corresponds to the case of injecting long-lived CBR traffic into the network. Our scheme consumes less energy than GAF even in this scenario for two reasons: (1) Grid leader election consumes a significant amount of energy, and (2) There still exists some routing redundancy in GAF due to the limitation on grid size that any two nodes in neighboring grids should be able to communicate directly.

In our scheme, a routing backbone backbone is not explicitly elected. Instead, the route discovery phase "automatically" selects the nodes to be turned on. Given a snapshot of the network at any particular time, the number of nodes that are active in our scheme is less.

Figure 11 also shows the "net" energy consumption of our prototype after eliminating that of beacon messages. We observe that a significant amount of energy is consumed by

beacon messages and so claim that it is necessary to devise more light-weight synchronization mechanisms to wake up nodes.

In summary, our prototype can achieve more energy conservation when compared to GAF by (1) adapting to traffic load and adjusting a node's duty cycle, and (2) invoking fewer intermediate nodes for data delivery. It should be noted that this comparison is biased against our prototype since in GAF, data sources/sinks need to stay awake all the time, while in our prototype, the data sources/sinks can go to sleep during idle periods.

VI. CONCLUSIONS AND FUTURE DIRECTIONS

In this paper, we present an on-demand power management framework that reduces energy consumption in ad hoc networks while maintaining effective throughput. It explores the design space between proactive and reactive power management approaches by adapting to the traffic load in ad hoc networks. In our framework, transitions between power management modes of a node are triggered by packet arrivals and the expiration of a soft state timer called the keep-alive timer. Various messages can serve to set up and refresh the keep-alive timer on demand.

We implement a prototype of our framework based on the IEEE 802.11 MAC in the ns-2 simulator. Simulation studies show that our framework consumes significantly less energy than a network without power management under a wide range of traffic patterns and mobility scenarios, while

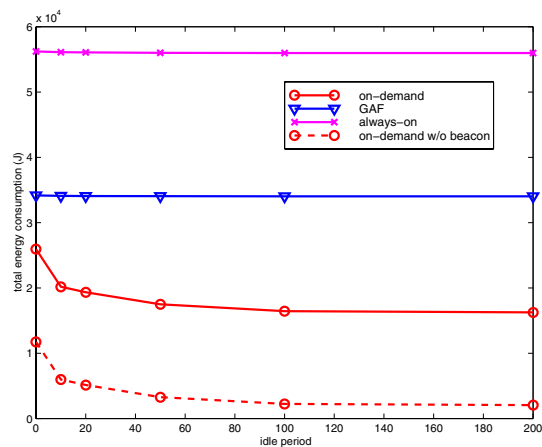


Fig. 11. Comparison with GAF, five on-off conns, on = 10s, 85 nodes, 1000mx1000m

maintaining a good balance between energy conservation and communication efficiency. In addition, the performance of our scheme degrades gracefully in the presence of high mobility. Comparisons with GAF show that our scheme can achieve better energy savings by adjusting a node's duty cycle based on its traffic load and invoking fewer intermediate nodes for data delivery.

Based on our current design and the results of the simulations presented in this paper, we plan to investigate the extension of our framework in the following directions.

Online adaptation of the keep-alive timers: From the simulation results, fixed keep-alive timers can sometimes be wasteful in the case of on-off traffic sources. We are currently working on the theoretical analysis of the trade-off between energy, latency and throughput. This will provide a basis for adapting the keep-alive timer values to the arrival pattern in the network.

Better handling of mobility: When mobility is high, intermediate nodes may remain in active mode longer than necessary, which will result in reduced energy savings. In addition, frequent broadcast messages for route discovery may hurt overall performance since broadcast messages cause all nodes to stay awake during a beacon interval. Techniques such as mobility prediction or proactive handoffs may be used to reduce the number of unnecessarily nodes that remain on despite route changes [18]. For on-demand routing protocols such as DSR that maintain routing caches, we are investigating the integration of power management with the caching strategy to better handle mobility.

Load balancing: To improve the longevity of the network, load balancing needs to be considered. In both Span and GAF,

load balancing is done by rotating the role of coordinators and grid leaders among neighboring nodes in a periodic fashion. Load balancing of this kind only addresses part of the problem since members of the CDS not involved in data transmission still need to remain on. So far, we have not considered load balancing issues in our framework. The correct policy for load balancing is dependent on the communication goals of the network.

REFERENCES

- [1] 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*, April 2001.
- [2] B. Chen, K. Jamieson, H. Balakrishnan, and R. Morris, "Span: An energy-efficient coordination algorithm for topology maintenance in ad hoc wireless networks," in *Proc. of ACM/IEEE 7th Int'l Conf. on Mobile Computing and Networking (MobiCom 2001)*, July 2001.
- [3] Y. Xu, J. Heidemann, and D. Estrin, "Geography-informed energy conservation for ad hoc routing," in *Proc. of ACM/IEEE 7th Int'l Conf. on Mobile Computing and Networking (MobiCom 2001)*, July 2001.
- [4] R. Kravets and P. Krishnan, "Application-driven power management for mobile communication," *Wireless Networks*, vol. 6, no. 4, 2000.
- [5] M. Stemm and R. Katz, "Reducing power consumption of network interfaces in hand-held devices," in *Proceedings of the 3rd International Workshop on Mobile Multimedia Communications*, 1996.
- [6] D. B. Johnson and D. A. Maltz, "Dynamic source routing in ad hoc wireless networks." draft-ietf-manet-dsr-04.txt, 2001.
- [7] UCB/LBNL/VINT, "Network simulator - ns (version 2)." <http://www-mash.cs.berkeley.edu/ns/>, January 2000.
- [8] R. Wattenhofer, L. Li, P. Bahl, and Y.-M. Wang, "Distributed topology control for power efficient operation in multihop wireless ad hoc networks," in *Proc. of IEEE INFOCOM*, 2001.
- [9] S. Narayanaswamy, V. Kawadia, R. S. Sreenivas, and P. R. Kumar, "Power control in ad hoc networks : Theory, architecture, algorithm and implementation of the COMPOW protocol," in *European Wireless Conference*, 2002.
- [10] J. P. Monks, V. Bharghavan, and W. Hwu, "A power controlled multiple access protocol for wireless packet networks," in *Proc. of IEEE INFOCOM*, 2001.
- [11] S. Singh, M. Woo, and C. S. Raghavendra, "Pamas: Power aware multi-access protocol with signalling for ad hoc network," *ACM Computer Communication Review*, pp. 5-26, July 1998.
- [12] C. Schurgers, V. Tsiatsis, S. Ganeriwal, and M. Srivastava, "Topology management for sensor networks: Exploiting latency and density," in *Proceedings of The Third ACM International Symposium on Mobile Ad Hoc Networking and Computing*, 2002.
- [13] W. Ye, J. Heidemann, and D. Estrin, "An energy-efficient MAC protocol for wireless sensor networks," in *Proceedings of the The 21st Annual Joint Conference of the IEEE Computer and Communications Societies*, 2002.
- [14] Y.-C. Tseng, C.-S. Hsu, and T.-Y. Hsieh, "Power-saving protocols for ieee 802.11-based multi-hop ad hoc networks," in *Proceedings of the The 21st Annual Joint Conference of the IEEE Computer and Communications Societies*, 2002.
- [15] C. E. Perkins, E. M. Royer, and S. Das, "Ad hoc on demand distance vector (AODV) routing." draft-ietf-manet-aodv-07.txt, 2000.
- [16] R. Zheng and R. Kravets, "On-demand power management for ad hoc network." UIUC Dept. of CS Technical Report No. UIUCDCS-R-2002-2280, 2002.
- [17] "CMU monarch extensions to ns." <http://www.monarch.cs.cmu.edu/>.
- [18] T. Goff, N. Abu-Ghazaleh, D. Phatak, and R. Kahvecioglu, "The broadcast storm problem in a mobile ad hoc network," in *Proceedings of the 7th Int'l ACM/IEEE Conference on Mobile Computing and Networking*, 2001.