Reducing Overhearing Energy in 802.11 Networks by Low-power Interface Idling

Subir Biswas
Electrical and Computer Eng. Department
Michigan State University
East Lansing, MI 48824
sbiswas@egr.msu.edu

Samir Datta
WinLab and Computer Sc. Department
Rutgers University
New Brunswick, NJ 08854
sdatta@paul.rutgers.edu

Abstract

In this paper we propose and analyze a new interface idling mechanism for improving energy efficiency of IEEE 802.11 based MAC hardware. A novel protocol state analysis technique is developed for detecting time windows during which a wireless interface consumes energy due to 802.11 overhearing. During this window, energy savings at the MAC layer is accomplished by forcing the wireless interface to a relatively lower-energy idling state. At the end of this window, the interface is transitioned back to its regular receiving mode. Energy savings are realized by exploiting the difference in power consumption between the overhearing state and the idling state. We evaluate the proposed protocol using NS-2 simulator. Simulation experiments demonstrate that the proposed mechanism is capable of reducing overhearing expenditure by up to 23% and can extend network life by up to 86%. Results also show that the proposed MAC layer idling is fairly insensitive to network loading.

1 Introduction

A mobile Ad-Hoc network (MANET) [1] is a group of mobile wireless nodes which, upon deployment, cooperatively form an infrastructure-less network without any centralized control and service infrastructure. Since the MANET nodes are typically run from limited energy portable batteries, a critical design issue for future wireless Ad-Hoc networks is the development of suitable communication architectures, protocols and services that reduce power consumption, thereby increasing the operational lifespan of network enabled wireless devices. Energy conservation in a MANET node not only maximizes its own operational lifespan but it can also help maximizing the network lifespan and deferring network partitioning.

In networks running IEEE 802.11 at MAC layer the expended *transmission energy* at a node is generally a tiny fraction of its total communication related energy expenditure. In addition to transmission and reception, there are three sources of non-essential energy expenditure [2] in an 802.11 network. The first one is due to *idling*, which corresponds to energy consumption when a wireless interface is neither transmitting nor receiving. In 802.11, even in idle state an interface must be up and ready to receive possible traffic. For older generation hardware, the consumption during this state is significant. Lucent's 915 MHz WaveLAN card, for instance, consumes 1.15W when idling, 1.2W while receiving and 1.6W while transmitting [3].

The second component is *overhearing*, during when an interface receives data and control packets that were transmitted to some other node. Overhearing is caused by the fact that when a unicast packet transaction is carried out in a node's immediate neighborhood, it does not have any mechanism for not to receive that packet. Energy expenditure during overhearing is same as that during reception.

The third component is *Erroneous Carrier Sensing (ECS)*, which corresponds to reception of signal with unacceptable Signal to Interference Ratio (SIR) [4]. In order to reduce collisions, 802.11 MAC protocol employs an extended carrier sense range which relies on a higher carrier sense sensitivity compared to the receive sensitivity. Although a higher carrier sense range reduces the collision rate, it increases the duration of ECS. ECS consumes same amount of energy as reception.

1.1 Characterization of Non-essential Energy Expenditures

To investigate the breakdown of non-essential energy components we conducted simulation experiments with NS-2 [5] for a typical MANET running 802.11b at the MAC layer and routing protocol AODV [6] at the routing layer. For the energy model, we have used Socket Communication Inc.'s Low Power Wireless LAN

Compact Flash Card that consumes 594mW while receiving and 924mW while transmitting at full power [7]. In addition, the card supports a low-power *idling* mode[1] in which it listens for carrier and consumes only 66mW. As soon as carrier is detected in the channel, the interface transitions into the receiving mode. A summary of all its interface states and their energy expenditure is presented in **Error! Reference source not found.**.

Experiments were performed with three traffic loading scenarios. For each scenario, the energy statistics was captured at the time of battery exhaustion for the node that was exhausted first. Percentage energy breakdown for transmission, reception and the non-essential expenditures are presented in **Error! Reference source not found.**. Following observations are to be made.

1. Transmission accounts for only a tiny fraction (3.4% to 6.8%) of the total energy expenditure
2. Non-essential expenditures (idling, overhearing and ECS) are responsible for the majority of consumption (88.7% for low load and 81.8% for high load)
3. All the energy components, except idling, increase with load; this is because with increasing traffic load part of the idling duration is converted to transmit, receive, overhearing and ECS.

1.2 Our Approach: Forced Idling

Goal of this work is to develop an energy conserving mechanism at MAC layer for limiting expenditure due to overhearing. Putting an interface to a low-power *sleep* mode [2,9] during overhearing is an obvious way to deal with it. However, the problem with forcing a wireless interface to a complete sleep state is that most of today's wireless interfaces typically take 10s of milliseconds to transition from sleep to any other state. Lucent's WaveLan card, for instance, takes around 100ms [3] for this transition. When compared to typical packet duration of few milliseconds, this transition latency is way too high.

The key idea in this paper is to limit overhearing expenditure by leveraging advanced power management features offered by next generation wireless hardware such as Intersil's PRISM GT 802.11 chipset [8] and Wireless LAN Compact Flash card from Socket Communication Inc [7].

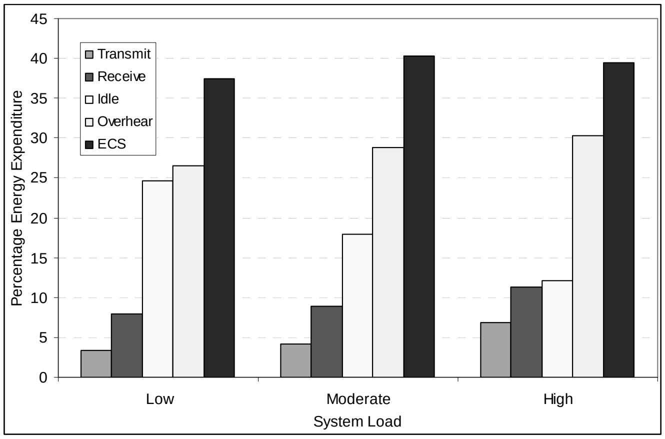

Figure 1: Breakdown of transmit, receive and non-essential energy expenditure in an 802.11 network

To avoid high wake-up latency from complete sleep, we propose forcing the interface to a low-energy *idling* state (e.g. the idle state as specified for the Socket Communication's low power wireless LAN card [7] in Table 1) during expected overhearing. In the *idling* state, instead of shutting the entire card off, certain parts of its electronics are kept powered so that the transition time to a non-idling state is only around 20µs. Our approach is to convert much of the overhearing time to the forced-idling. Due to the low idle consumption numbers (66mW compared to 594mW for overhearing) we achieve overall reduction in consumption due to overhearing.

The main contribution of this paper is an 802.11 protocol extension for limiting overhearing energy by forcing an interface to idle based on the following approach. While in overhearing state, if an interface detects a packet transaction in its neighborhood, it transitions to a lower energy idle state till the packet transaction is over. This way, it can reduce the energy consumed due to overhearing. After the surrounding transaction is over, the interface transitions back to the overhearing state. We implement this mechanism and show that this forced state-switching can extend network life by reducing the overhearing expenditure.

Note that the mechanisms for reducing erroneous carrier sensing (ECS) expenditure are outside the scope of this paper and they are reported in a separate publication [10].

---

[1] A number of next generation wireless chipsets support advanced power management modes. Another notable example is GlobaspanVirata's PRISM GT (formerly from Intersil) 802.11 chipset [8] that supports four low-power modes with current consumption ranging from 30mA to 190mA.

The rest of the paper is organized as follows. Section 2 provides MAC layer protocol details for energy savings on 802.11 overhearing. Details about the simulation experiments and performance results are furnished in Section 3. Finally, we summarize the results and conclude the paper in Section 4.

Interface State	Description	Consumption
Idling	Interface is ready to receive but no carrier is present	66mW
Receiving	Receiving packets destined to itself	594mW
Overhearing	Receiving packets not destined to itself	594mW
Erroneous Carrier Sensing (ECS)	Receiving signal with unacceptable SIR	594mW
Transmitting	Transmitting packets	924mW max

Table 1: Energy consumption of Socket Communication Inc.'s Low Power Wireless LAN Compact Flash card

2 O-Idling: MAC Layer Idling for Overhearing Reduction

Overhearing at a MANET node is caused by reception of packets that are not meant for it. For example in the liner topology in , when node B sends a data packet to node C, the packet is overheard by node A. Our goal is to avoid spending energy on overhearing by forcing A's radio interface to transition to a low-energy idling mode during B's transmission to C.

We exploit the NAV (network allocation vector) mechanism in 802.11 to accomplish forced idling. When A overhears the B→C RTS, it looks at the NAV value in RTS which indicates the duration of packet transfer till the end of C→B ACK transmission. Upon overhearing the RTS, A should force its radio interface to idling state and schedule a transition back after the end of the C→B ACK transmission. During this period (marked as "o-idling" for node A in), node A consumes energy at the idling rate compared to the higher overhearing rate, and that translates to savings.

Similar techniques are applicable for node D, which forces its interface to idle based on the NAV it has found in the overheard C→B CTS. Since D doe not overhear B→C data, even without forced idling it is likely to spend a significant part of that NAV duration idling, except when it overhears C→B ACK and possible erroneous carrier sense from nodes that are far. As a result, reduction of overhearing energy at node D is less than that at node A. At D, however, a certain amount of erroneous carrier sensing energy is saved due to o-idling.

Now let us consider the situation in which an X→Y packet transmission is initiated while node A was in o-idling state. Since A has missed the Y→X CTS, when it comes out of the o-idling it has no knowledge of the ongoing X→Y data transmission. At this point if A has a packet to transmit to either Y or B, it may potentially collide with X→Y data at node Y. We call this an *idling-collision*. Similar issues exist after node D comes out of its o-idling. To address this, we impose a restriction that after the end of o-idling, a node should wait for duration T_{guard} before attempting a transmission. In our example, if X→Y transmission and its ACK end before A's T_{guard} is over then node Y will be protected from the idling-collisions. Large T_{guard} values provide better protection from idling-collisions, but it also brings the channel capacity down by forcing nodes not to transmit during that guard period. On the other hand, for very small T_{guard} values, the effective channel capacity will suffer due to idling-collisions. Therefore, an optimum T_{guard} has to be set.

The most conservative value of T_{guard} should be set as $T_{guard}^{\max} = T_{data}^{\max} + SIFS + T_{ack}$, where $T_{data}^{\max}$ is the maximum data packet duration, $SIFS$ is the short inter-frame spacing in 802.11 and T_{ack} is the ACK packet duration. With this guard time, the collision between X→Y data and A→Y transmission can be avoided in the very unlikely event when Y→X CTS arrives at node A just before it comes out of the o-idling state. Since this is less likely, in practice a smaller guard time is desirable. It is more so because often times we observed that in Ad-Hoc topologies when a node such as A receives RTS from B, its immediate neighbors such as Y also receives the same RTS. As a result, just like A, node Y is also forced to the o-idling state and that would prevent an idling-collision as described above. Through experiments we found that a value T_{guard} around $T_{guard}^{\max}/16$ provides maximum overhearing energy savings for the topologies and scenarios that we have experimented with.

To summarize, whenever a node overhears an RTS or CTS control packet that is not destined to it, it forces its interface to transition to a low energy idling state. The node also sets a timer based on the NAV value in the received control packet. Upon expiration of that timer, the node takes the o-idling condition off and makes the interface wait for T_{guard} duration, before it can transmit any control packet. During this guard time a node can be in overhearing, erroneous carrier sensing or idling state depending on its surrounding channel activities.

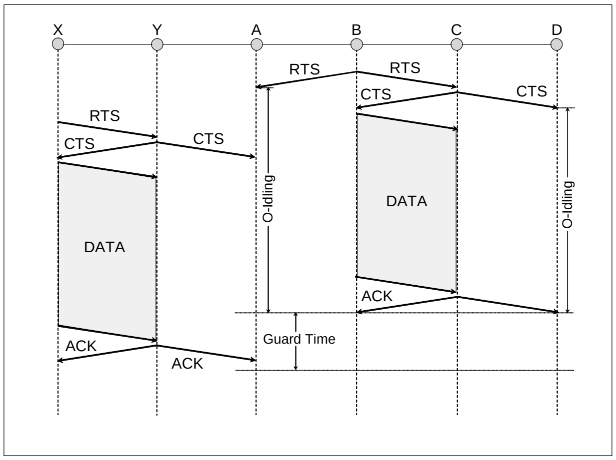

Figure 2: Mechanism and timing for *o-idling*

O-idling has the following shortcomings. First, it cannot eliminate overhearing completely because of the guard period during which a node may have to overhear surrounding activities. Second, the forced inactivity during the guard period can result into reduced throughput as well as increased packet delivery delay. By proper dimensioning of T_{guard} both these effects can be mitigated to some extent. There is another source of throughput reduction and that is due to unsuccessful RTS and CTS. If a node is forced to o-idle after overhearing an RTS or CTS which do not culminate to a successful data transmission then o-idling proves to be redundant and reduces the nodes transmission time. According to our observation however, this effect is not very pronounced. Finally, the o-idling mechanism, as described above, does not work when 802.11 fragmentations [13] are turned on. That is because with fragmentation option turned on, while transmitting a data segment, the transmitting node can extend the current NAV if there is at least one more fragment to transmit. Once a node's interface has been forced to o-idling, it does not have any way of knowing about that extended NAV. As a result, the node will end its o-idling sooner than necessary and thus not saving as much energy as it possibly could. To mitigate this problem somewhat, the node can look at the NAV in overheard data packets and go back to o-idling repeatedly.

Two other papers attempt to achieve similar objectives, but through different means. In [11], overhearing is controlled using forced sleeping when there are ongoing packet transactions in a nodes neighborhood. The main difference of our approach with theirs is that they rely on an out-of-band busy-tone signaling channel to exchange information about packet transmission schedules. Whereas in our approach, the neighborhood transmission state is inferred from the 802.11 Network Allocation Vectors (NAV) and that way the need for two radio transceivers is avoided. Also, as done in [11], turning an interface off completely has the problem of unacceptable transition latency as described in Section 1.2. To avoid this problem, we force the interface to a low energy idling state but not to a completely off state such as sleeping. While the energy savings are relatively lesser in our approach, it is more feasible from an implementation standpoint.

The authors in [2] propose a mechanism in which an interface is tuned off based on its internal data activity. It implements a MAC layer state machine that monitors traffic arrival from upper layers for transmission and turns the radio-interface off if there is data inactivity longer than a pre-defined period. The interface is again turned on after it spends a pre-defined time in the off state. The main difference between this approach and ours is that in ours a node is forced to a low energy state based on the transmission activity in its neighborhood, whereas in [2] it is done purely based on the local data activity within the node. The trade-off in [2] is that a node can sleep longer and save more energy while the lack of synchronization between a transmitter and a receiver may lead to packet drops. This will happen when a transmitter sends RTS and its intended neighbor receiver is in a sleep state and cannot reply to that RTS. If this happens more than retransmission threshold times then the packet is dropped. To address this problem within routing protocols such as AODV [6], [2] proposes a RREQ blanket retransmission mechanism to make sure those RREQ packets will eventually be delivered. This causes heavy control overhead and increased route establishment latency. Moreover, after a route is created, regular data activity on nodes along that route prevents them from sleeping altogether. As a result, they consume energy due to overhearing. In our approach, since a node switches between idling and other states based on its neighborhood transmission, it always saves overhearing energy irrespective of its internal data activity.

3 Performance

Performance of O-idling has been evaluated using NS-2 simulation package. A 60-node network has been simulated within a 550mx120m rectangular area. AODV and 802.11 have been used as the base routing and MAC layer protocols. For mobility pattern, we assume that each node pauses for a fixed duration between movements and after each pause it moves towards a random direction with an average speed of 6.4m/s. Traffic is generated as 512 byte long CBR packets. Various inter-packet spacing, ranging from 0.3s to 1.2s, are used for varying the loading conditions. CBR sessions are generated between randomly chosen source-destination pairs with each session lasting for 500s. For energy model, we use the power consumption figures for Socket Communication's low power wireless LAN card as summarized in **Error! Reference source not found.**.

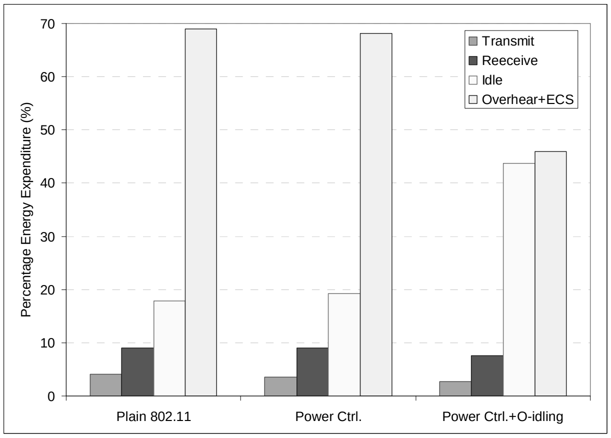

Figure 3: Impacts of power control and O-idling on non-essential energy expenditures

All our experiments were run in the presence of Transmission Power Control (TCP) [12]. We implement an open-loop TCP. Each node maintains a neighbor table that is built using AODV hello mechanism. For each neighbor, its geographical coordinate[2] is stored for computing the physical distance. A node sends RTS, CTS and the broadcast packets with maximum available transmission power. However, for the data packets, a sender evaluates the physical distance to the receiver [12] and computes the necessary transmission power based on the propagation model and antenna properties. A two-ray-ground model has been used for our implementation.

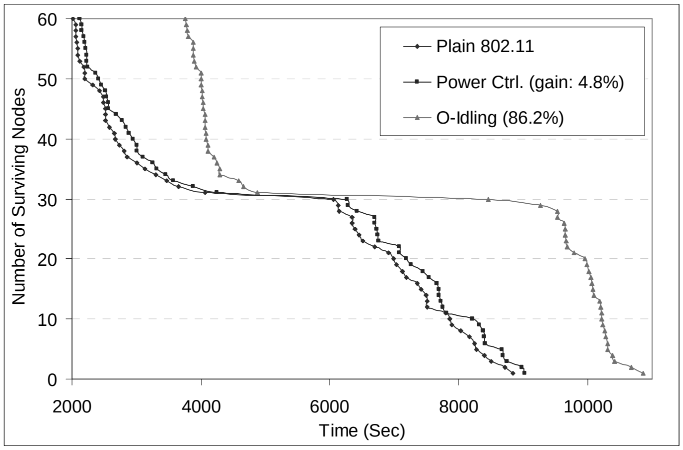

Figure 4: Impacts of forced idling on network life

Our first experiment is to characterize the energy impacts of O-idling and transmission power control, and to compare them with the base case of unmodified 802.11 with AODV routing. The base case is marked as "Plain 802.11" in **Figure 3**. As a first step we apply transmission power control which reduces the transmission expenditure somewhat (from 4.8% to 3.5%) but as expected, it does not have any appreciable impact on overhearing. However, applying O-idling with power control reduces the overhearing energy substantially by converting most of the overhearing duration to low-energy idling. Combined overhearing and ECS expenditure is reduced from 69% to 45.9% and the reduction is solely due to the conversion of overhearing time to idling time. As a result, O-idling causes a significant increase in idling (from 18% to 43.7%). However, due to almost an order of magnitude consumption difference (66mW during idling as opposed to 594mW during overhearing) the overall consumption is reduced.

The effects of energy savings on overall network life are shown in **Figure 4**. The graphs show number of surviving nodes as the time progresses. For a given scenario, the measure of goodness for a energy-

---

[2] We assume that the nodes are equipped with outdoor GPS or low-cost indoor Local Positioning System (LPS).

saving protocol is indicated by how far right the graph is pushed from the base case (marked as “plain 802.11”) scenario. For all experiments, out of 60 nodes 30 are given an initial energy of 500Joule and the rest 30 are given 1000Joule. The low and high energy nodes are chosen randomly for energy heterogeneity. It is due to this differential energy distribution, we see that each graph in **Figure 4** has an upper half that corresponds to the battery exhaustion of 30 low-energy nodes and then the lower half that corresponds to high-energy nodes.

We define the lifespan gain as the percentage increase of the life of the first battery-exhausted node. While transmission power control provides a gain of modest 4.5%, O-idling, combined with power control, offers an 86% gain over the base case.

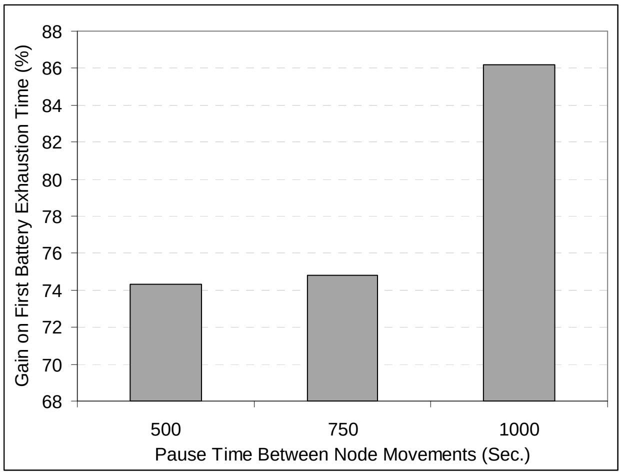

Figure 5: Effects of node mobility on network life

Impact of node mobility on network lifespan gain is illustrated in **Figure 5**. Change of mobility is accomplished by adjusting pause times between node movements. Changing pause time from 500s to 1000s improves the gain in lifespan from approximately 74% to 86%. Errors in idling-duration estimations (see Figure 2), due to mobility, account for this loss of lifespan extension gain with higher mobility.

Packet drop rate in **Figure 6** is defined as the number of packets dropped for each successfully delivered packet. Packet drops are contributed by MAC layer buffer overflow and drops after pre-defined number of MAC retransmissions, which happens during local congestions. In this experiment we study overall packet drop rates for the base case as well after applying power control and O-idling. As evident from the diagram, application of the MAC layer idling mechanism raises the drop rate from around 0.6% to 4.2%.

Note that the idling durations in O-idling (see) are only estimations and there are estimation errors, which can cause over-idling. In the events of over-idling a node may potentially loose packets that are transmitted towards it. These losses account for the increase in packet drop rate.

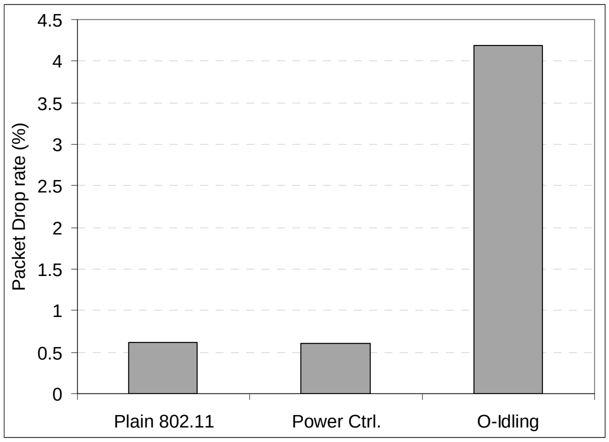

Figure 6: Impacts of forced idling on packet drop rate

Increased drop rate is the price being paid for improved energy efficiency. With proposed MAC layer idling mechanism we observe a lifespan extension of 86% over the base case. This gain is traded for an increased packet drop rate from 0.6% to 4.2%. This would impact data oriented applications such as TCP more than the stream oriented applications that can tolerate higher packet drop rates. We did not observe any noticeable impact of the proposed energy saving mechanisms on packet delivery hop count and its associated delays. Results also show that the network lifespan extension is fairly insensitive to network loading.

4 Summary and Conclusions

In this paper we have proposed and analyzed a forced interface idling mechanism for improving energy efficiency of IEEE 802.11 based MAC hardware. A novel protocol state analysis technique is developed for detecting time windows during which a node consumes energy due to overhearing. During this window, energy savings at the MAC layer is accomplished by forcing a wireless interface to a relatively low-energy idling state. Experimental results demonstrate that the proposed mechanism is capable of reducing overhearing expenditure by up to 23% and can extend network life by up to 86%. This significant gain in network lifespan is traded for increased packet drop rates. The maximum packet drop rate that we observed was 4.2% with an overall network life extension of 86%. Results also show that the proposed MAC layer idling is fairly insensitive to traffic loading. To summarize, our proposed MAC idling mechanisms offer an efficient means for cutting energy expenditures due to 802.11 overhearing. This is accomplished by leveraging the low-power idling modes that are supported by next generation wireless interface cards.

5 Reference

1. J. Macker and S. Corson, "*Mobile ad hoc networks(MANET)*," http://www.ietf.org/html.charters/manet-charter.html, 1997. IETF Working Group Charter

2. Ya Xu, John Heidemann, and Deborah Estrin, "*Adaptive energy-conserving routing for multihop ad hoc networks,*" Research Report 527 USC/Information Sciences Institute, October 2000.

3. M. Stemm and R. Katz, "*Measuring and Reducing Energy Consumption of Network Interfaces in Hand-held Devices,*" IEICE Transaction on Communications, vol. E80-B, no. 8, August 1997

4. M. Takai, J. Martin, and R. Bagrodia, "*Effects of Wireless Physical Layer Modeling in Mobile Ad Hoc Networks,*" In Proceedings of MobiHooc2001, Long Beach, CA, October 2001

5. ns-2 Manual, http://www.isi.edu/nsnam/ns/doc/index.html

6. C. Perkins, E. Royer, and S. Das, "*Ad hoc on demand distance vector (AODV) Routing,*" IETF Internet Draft, http://www.ietf.org/internet-drafts/ draft-ietf-manet-aodv-07.txt, Nov 2000

7. Product description of SOCKET Communications Inc., low power wireless LAN card http://www.socketcom.com/pdf/wlan/wlandatasheet.pdf

8. Product description of PRISM GT WLAN solution, http://www.globespanvirata.com/prismgt.html

9. N. Gupta and S. Das, "*Energy-aware On-demand Routing for Mobile Ad-Hoc Networks,*" In Proc. Of IWDC 2002, Calcutta, India, December 2002

10. S. Datta and S. Biswas, "Energy Conservation by Cross-layer Protocol Design in Wireless Ad-Hoc Networks," Submitted to MobiHoc 2004

11. S. Singh and C.S. Raghavendra, "*PAMAS: Power Aware Multi-access Protocol with Signaling for Ad Hoc Networks,*" ACM Computer Communication Review, vol. 28, no. 3, July 1998

12. D. Qiao, S. Choi, A. Jain and K. G. Shin, "*Adaptive Transmit Power Control in IEEE 802.11a Wireless LANs,*" In proc. Of IEEE VTC Spring, Korea, April 2003

13. B. O'Hara and A. Petrick, "*802.11 Handbook: A Designer's Companion,*" Standards Information Network, IEEE Press, 1999.