

Investigating the Energy Consumption of a Wireless Network Interface in an Ad Hoc Networking Environment

Laura Marie Feeney, Martin Nilsson

Abstract—Energy-aware design and evaluation of network protocols requires knowledge of the energy consumption behavior of actual wireless interfaces. But little practical information is available about the energy consumption behavior of well-known wireless network interfaces and device specifications do not provide information in a form that is helpful to protocol developers. This paper describes a series of experiments which obtained detailed measurements of the energy consumption of an IEEE 802.11 wireless network interface operating in an ad hoc networking environment. The data is presented as a collection of linear equations for calculating the energy consumed in sending, receiving and discarding broadcast and point-to-point data packets of various sizes. Some implications for protocol design and evaluation in ad hoc networks are discussed.

Keywords—energy consumption, IEEE 802.11, ad hoc networks

I. INTRODUCTION

Because mobile devices are dependent on battery power, it is important to minimize their energy consumption. The energy consumption of the network interface can be significant, especially for smaller devices. Most research in energy conservation strategies has targeted wireless networks that are structured around base stations and centralized servers, which do not have the limitations associated with small, portable devices.

By contrast, an ad hoc network is a group of mobile, wireless hosts which cooperatively form a network independently of any fixed infrastructure. The multi-hop routing problem in ad hoc networks has been widely studied in terms of bandwidth utilization, but energy consumption has received less attention. It is sometimes (incorrectly) assumed that bandwidth utilization and energy consumption are roughly synonymous. Recently, there has been some study of energy-aware ad hoc routing protocols, particularly for distributed sensor networks. In this context, energy is often treated as an abstract “commodity” for purposes of minimizing cost or maximizing time to network partition.

We believe that energy-aware design and evaluation of network protocols for the ad hoc networking environment requires practical knowledge of the energy consumption behavior of actual wireless devices. In addition, it is important to present this information in a form that is useful to protocol developers: the total energy costs associated with a packet containing some number of bytes of data. Device specifications, which indicate the current draw while transmitting and receiving, are somewhat unhelpful in this respect.

This paper describes a series of experiments which obtained detailed measurements of the energy consumption of a Lucent WaveLAN IEEE 802.11 wireless network interface operating in ad hoc mode. The data is presented as a collection of linear equations for calculating the energy consumed in sending, re-

ceiving and discarding broadcast and point-to-point data packets of various sizes. While not intended to account for all IEEE 802.11-based products or for all possible factors affecting energy consumption at the wireless interface, the results provide a solid experimental basis for energy-aware design and evaluation of network-layer protocols operating in the IEEE 802.11 environment.

The results suggest new perspectives on design of ad hoc networking protocols:

First, it is clear that energy consumption and bandwidth utilization are not synonymous. It is necessary to consider not only the cost of transmitting a packet, but also of receiving, and even of discarding it. Protocol designers must therefore consider the proportions of broadcast and point-to-point traffic used by the protocol. Because channel acquisition overhead is large, small packets have disproportionately high energy costs. Promiscuous mode operation, which is irrelevant to bandwidth utilization, also incurs some energy cost.

Second, the relationship between transmit speed and overall energy consumption is complex. Reduced data transmit and receive times have only limited impact on per-packet energy consumption, due to the high fixed overhead. There are other trade-offs involved: the reduced transmission range associated with higher data rates both increases the number of hops required in a multi-hop routing environment, but decreases the number of neighbors affected by each transmission.

Third, ad hoc mode operation incurs an extremely high idle cost compared to operation in conjunction with a base station. Routing protocols which are structured to emulate some of the energy efficiency associated with BSS mode operation should be investigated.

II. RELATED WORK

Experimental Results

There are few published measurements of the energy consumption of network interfaces. In particular, there are no detailed measurements of the per-packet energy consumption of an IEEE 802.11 network interface operating in ad hoc mode. In general, device specifications, which show current draw while transmitting and receiving aren't sufficient to calculate per-packet energy consumption. For example, idle mode, switching from transmit to receive mode and the effects of internal energy management strategies may not be reflected.

Experiments measuring the power consumption of IEEE 802.11 PC cards are reported in [10]. Emphasizing operation in conjunction with a base station, the methodology is based on sampling the current draw of an otherwise idle laptop as it sends

and receives traffic over an extended period of time. By comparing the total energy used while sending or receiving traffic with the energy consumed by an idle laptop, the total cost of processing traffic can be computed. This approach has the advantage of measuring total system cost, but provides little packet-oriented detail. Also, the results are potentially more dependent on particular system-wide energy management techniques than on the behavior of the network interface.

Packet-oriented measurements of the energy consumption of wireless interfaces are described in [7] and the experiments described below are very closely based on this methodology. However, this work only includes measurements of the pre-IEEE 802.11 WaveLAN interface (and several other devices). In addition to updating these results to include more recent hardware, the current work examines ad hoc mode operation of the network interface and complexities introduced by the IEEE 802.11 protocol. It is also unique in investigating the behavior of non-destination hosts which overhear wireless traffic.

Energy-aware Protocol Design and Evaluation

A detailed analytical study of the energy efficiency of a number of MAC-layer protocols, including IEEE 802.11, is presented in [2]. This probabilistic analysis examines the effectiveness of various media acquisition strategies in the presence of contention. Contention is dependent on many factors, such as communication patterns, node density and RF transmission characteristics, that are difficult to reproduce experimentally. We therefore attempted to minimize the possibility of media contention and retransmissions in our measurements.

Most energy-conserving link-layer protocols are targeted toward the centralized, base station environment. Such protocols usually rely on a resource-rich base station to moderate communication among hosts, scheduling and buffering traffic to reduce contention and allowing the more limited mobile devices to spend as much time as possible in a low-power consumption sleep state. Unfortunately, these strategies have limited applicability in the ad hoc environment, in which there are no fixed base stations and mobile hosts may have limited buffering capability and unpredictable connectivity.

Evaluating the energy efficiency of network-layer protocols has proven to be a surprisingly subtle task. Energy consumption and bandwidth utilization are substantively different metrics: the former must reflect costs for sending, receiving and discarding traffic and emphasizes the differences between broadcast and point-to-point traffic. A simulation study [4] of the energy consumption of two well-known ad hoc routing protocols running over IEEE 802.11 demonstrated that an energy-oriented performance evaluation may come to quite different conclusions than a bandwidth-oriented one. This result, which was based in part on earlier versions of these experiments [6], has provided impetus for the more comprehensive investigation presented here.

Application-level energy conservation strategies [7], [10] take advantage of usage patterns associated with activities such as email retrieval and web browsing. This allows the network interface to spend as much time as possible in an inactive, reduced power consumption state with minimal impact on the performance perceived by the user. Such techniques are not applicable

to an ad hoc network. Because the hosts in an ad hoc network also form its routing infrastructure, it is difficult or impossible to predict when it is "safe" for a network interface to enter a low-power sleep state. Even entirely quiescent nodes may be (or become) essential to maintaining network connectivity or providing other essential network services.

There has also been some recent interest in energy-based techniques for ad hoc sensor networks, in which a collection of specialized sensors cooperatively forward sampled data to more powerful hosts for further processing or other action. Routing protocols which seek to maximize the connected lifetime of the network by energy-aware load balancing are presented in [1], [8]. Methods for selecting transmit power levels to maximize network lifetime and reduce spatial interference are presented in [1], [13]. However, energy is often treated as abstract commodity and subtle issues such as those suggested above are not addressed.

III. OVERVIEW

The goal of this work was to investigate the energy consumption of a wireless network interface via direct measurement

Is the network interface a significant factor in overall system energy consumption?

A variety of measurements of the power consumption of an idle laptop computer are found in [3], [10], [7]: reported results range from 6 to 14 W. However, the growth of mobile computing is leading to the development of low-power "mobile" processors and systems. Detailed measurements of the energy consumption of Compaq's experimental Itsy "pocket computer" are reported in [3]. This PDA has an idle power consumption of 100 - 200 mW; most applications exhibit peak power consumption of 1 - 1.5W. Transmeta's proposed "all-day mobile computer" [16] is expected to consume 5 - 6W, of which the active CPU will account for 1-2W.

Lucent IEEE 802.11 WaveLAN device specifications suggest transmit and receive power consumption of about 1.5W and 1W, respectively. A host generally spends only a small portion of its time sending and receiving traffic, so idle power consumption is important to overall energy costs. Operating in ad hoc mode, the idle power consumption is nearly as large as that of receiving data. Reducing the energy consumption of the network interface is therefore extremely important.

What effect does ad hoc operation have on power consumption?

IEEE 802.11 defines two primary modes of operation for a wireless network interface: base station (BSS) mode and ad hoc mode. Every mobile host operating in BSS mode must be in transmission range of one or more base stations, which are responsible for buffering and forwarding traffic between hosts. Hosts can send outgoing traffic to the base station anytime and periodically poll the base station to receive incoming traffic. The remainder of the time is spent in a sleep state, from which the interface must explicitly wake up in order to send or receive traffic. The base stations' guaranteed availability and buffering and traffic management capabilities are required to support this energy conserving functionality.

Ad hoc mode operation does not use any base station infrastructure: nodes communicate directly with all other nodes

that are in wireless transmission range. Because there are no base stations to moderate communication, hosts must always be ready to receive traffic from their neighbors. An network interface operating in ad hoc mode does not sleep; it has a constant idle power consumption which reflects the cost of listening to the wireless channel. This cost, which has been measured, but is not described in device specifications, is only slightly less than that of actually receiving traffic.

How can we model per-packet energy consumption?

In the simple case, the energy consumed by the network interface when a host sends, receives or discards a packet can be described using a linear equation

$$Energy = m \times size + b.$$

Trivially, there is a fixed component associated with device state changes and channel acquisition overhead and an incremental component which is proportional to the size of the packet. Experimental results confirm the accuracy of the linear model and are used to determine values for the linear coefficients m and b for various operations.

The model does not consider the case of link-layer fragmentation. It is expected that the linear model would continue to apply, with each instance of fragmentation introducing a small step discontinuity reflecting a fixed fragmentation overhead.

The model also does not consider energy consumed in unsuccessful attempts to acquire the channel (media contention), or in messages lost due to collision, bit error or loss of wireless connectivity. Such effects are difficult to obtain for controlled experimental measurements. While these phenomena are clearly important for the energy consumption behavior of a network, they are probably best examined probabilistically in the context of a specific model of host density, traffic load and wireless transmission environment.

What are the relative costs of sending, receiving and discarding traffic? What are the relative costs of large and small packets? Of broadcast and point-to-point traffic? Of promiscuous mode operation?

In contrast with bandwidth metrics, which count the number of packets or bytes sent over the wireless media, energy consumption metrics must account for the reaction of every network interface within wireless transmission range of the participating hosts.

The relative magnitudes of the various m and b coefficients also indicate the amount of per-packet energy consumption overhead.

A packet may be sent as broadcast or point-to-point traffic. The former is received by all hosts within transmission range; the latter is discarded by non-destination hosts, unless they are operating in promiscuous mode, while the MAC traffic is processed by all hosts in range of either the sender or the destination. In every case, energy will be consumed at the all relevant network interfaces. It is important to note that the costs of receiving and discarding are multiplied by the number of hosts which receive or discard the traffic. Energy consumption is affected by node density.

Protocols can be designed to use different combinations of large and small packets. Some use techniques which piggyback their data onto existing traffic. Protocols can also be designed to

use different combinations of broadcast and point-to-point messages. For the former, all nodes in wireless range of the sender must process the traffic. For the latter, non-destination hosts may discard traffic, unless they are operating in promiscuous mode. In addition, the collision avoidance and acknowledgment mechanisms used by the IEEE 802.11 protocol differ for broadcast and point-to-point messages. We note that bandwidth metrics do not address these issues.

IV. MEASURING ENERGY CONSUMPTION

The test cards were popular and widely available 2.4GHz DSSS Lucent IEEE 802.11 WaveLAN PC “Bronze” (2Mbps) and “Silver” (11Mbps) cards. The test host was an IBM ThinkPad 560, running FreeBSD 4.0 and the freely available WaveLAN IEEE 802.11 driver written by Bill Paul (wpaul@ctr.columbia.edu). The machine ran with power management turned off in order to avoid unexpected interaction with system facilities. UDP test traffic was generated at the rate of 10-20 packets per second on otherwise idle machines running in single-user mode.

Significant efforts were made to avoid interference which might lead to media contention or retransmissions. The tcpdump facility was used to ensure that no other traffic was present on the channel. Channels known to be in use by various nearby WLAN installations or which demonstrated unusually noisy idle power consumption were also avoided. It was probably impossible to avoid all sources of RF interference, including GSM and DECT phones, WLAN installations and various laboratory equipment. However, there was little evidence of significant retransmissions.

The circuit and methodology were closely based on the experiments reported in [7]. Energy consumption was determined by direct measurement of the input voltage and current draw at the network device. The latter was obtained by inserting a small resistance in series with the device.

The test circuit was built using a Sycard PCCextend 140A CardBus Extender [15]. The extender is like a breakout box: it is inserted into the PC card slot on the host and the card to be tested is inserted into the card connector on the extender. The V_{cc} line can be isolated: $0.5 \Omega \pm 2.5\%$ test resistance was inserted at this point¹. See Figure 1 for a diagram of the current measurement portion of the circuit. Data measurements were made using a Tektronix 100MHz digital oscilloscope and 15 MHz 1X probes. In this configuration, the voltages of interest can be measured with a scope accuracy of $\pm 5.2 - 6.8\%$.²

The input current, $i_{in}(t)$, was determined by measuring the voltage $v_r(t)$ across the test resistance, R . The input voltage, $v_{in}(t)$, is affected by fluctuations in V_{cc} as well as by the varying $v_r(t)$. For small R , the latter is small compared to V_{cc} ($< \pm 1V$). The input voltage is therefore approximated by a constant, V_{in} .

The instantaneous power consumption is the product of the

¹The extender actually allows all control and data lines to be examined using a scope or logic analyzer; it is usually used for testing and debugging PC cards.

²The scope traces reproduced in Section V were made using 50MHz Fluke oscilloscope and 1Ω resistance. The Fluke instrument is less sensitive, but has a more manageable downloading facility.




Fig. 1. Test circuit

input voltage and current:

$$P(t) = V_{in} \frac{v_r(t)}{R}.$$

Total energy consumed over an interval $t_0 \dots t_1$ is the integral of power consumption over time:

$$E_{t_0 \dots t_1} = \frac{V_{in}}{R} \int_{t_0}^{t_1} v_r(t) dt.$$

When operating in ad hoc mode, the idle power consumption of the network interface is constant. (This is confirmed experimentally in Section V.)

$$P_{idle} = \frac{V_{in}}{R} v_{r_{idle}}.$$

It is much easier to measure energy consumption over an arbitrary interval than it is to consistently identify the precise duration of a transmit or receive event. Instead, we measure the total energy consumed during an interval that roughly brackets the event of interest and subtract away the constant idle power consumption.

If $t_0 \dots t_1$ is such an interval, then the additional energy consumed processing a packet is given by:

$$\frac{V_{in}}{R} \overline{v_r}(t_1 - t_0) - P_{idle}(t_1 - t_0),$$

where $\overline{v_r}$ is the mean value of v_r in the interval $t_0 \dots t_1$. This is easily computed on-board modern digital oscilloscopes³.

This technique can be used to measure the additional energy consumed by a network interface as it sends, receives or discards broadcast and point-to-point messages and it is in this form that our results are presented. Coefficients for the equations defined by the linear formulation above can be determined by performing these measurements for packets of various sizes and applying linear regression.

V. VISUAL AND ANALYTIC OVERVIEW OF THE IEEE 802.11 PROTOCOL

Each section below describes part of the IEEE 802.11 protocol, defines coefficients for use in a linear formulation and

³It is not appropriate to make A/C measurements because the data traffic is a very low frequency signal.




Fig. 2. Sending a packet




Fig. 3. Sending 2Mbps broadcast UDP/IP traffic (256 bytes)

presents some representative oscilloscope traces of the measurements⁴. The bottom and right axes show the scope measurements (time and $v_r(t)$, respectively). The left axis shows the corresponding instantaneous power consumption.

Broadcast Traffic

Before sending broadcast traffic, the sender listens briefly to the channel. If no signal is detected, the message is sent. Otherwise, the sender must back off and retry later. As noted above, we do not model this case.

The network configuration for this case is shown in Figure 2. Using a linear equation, where m represents the incremental cost and b represents fixed costs:

$$E_{broadcast-send} = m_{send} \times size + b_{send(bcast)}$$

$$E_{broadcast-recv} = m_{recv} \times size + b_{recv(bcast)}$$

This behavior is seen in the oscilloscope traces shown in Figures 3 and 4. The most obvious feature is the extremely high idle mode power consumption in ad hoc mode: actually receiving data requires only marginally more energy than waiting for it.

In Figure 3, it is easy to examine the 2Mbps rating of the interface. The payload is 256 (228 (data) + 8 (UDP) + 20 (IP)) bytes; MAC overhead includes the 24-byte PLCP header (which is transmitted at 1Mbps) and the 20-byte MAC frame header. This implies a 1.3 ms data transmit time, which is about what we see in the trace.

⁴Note that these plots are not screenshots of the scope traces. They were obtained by downloading waveform coordinates from a Fluke ScopeMeter and re-plotting them using 'gnuplot'.




Fig. 4. Receiving 2Mbps broadcast UDP/IP traffic (256 bytes)

A host that is not in transmission range of the sender cannot detect the sender's signal when it senses the channel and will proceed to send its own transmission. Both signals will be received at any host that is in range of both senders. Depending on relative signal strengths at each receiver, one or both packets will be lost due to collision. This is known as the hidden terminal problem.

Point-to-point Traffic

While the IEEE 802.11 protocol does not support media reservation for broadcast traffic, point-to-point traffic can use collision avoidance techniques to reduce the impact of the hidden terminal problem.

Before sending a point-to-point transmission, the source broadcasts a RTS (request-to-send) control message, specifying a destination and data size (duration). The destination responds with a CTS (clear-to-send) message. If the source does not receive the CTS, it may retransmit the RTS message. On receiving the CTS, the source sends the DATA and awaits an ACK from the receiver.

Any host that hears the RTS/CTS exchange must refrain from transmitting for the specified duration. This "virtual carrier sense" reduces, but does not eliminate, the possibility of collision at the destination node.

The network configuration and linear formulation used for these measurements are analogous to the ones used for broadcast traffic.

$$E_{point-to-point-send} = m_{send} \times size + b_{send(p2p)}$$

$$E_{point-to-point-recv} = m_{recv} \times size + b_{recv(p2p)}$$

The differences between the values of b_{send} and b_{recv} for broadcast and point-to-point traffic reflect the difference between the two kinds of channel access. The incremental cost of sending or receiving data once the channel is acquired is expected to be the same both broadcast and point-to-point traffic.

The oscilloscope traces are shown in Figures 5 and 6 and the various elements of the media reservation protocol can be easily identified.

For small packets, this media reservation exchange represents considerable overhead. Therefore, for packets smaller than (a




Fig. 5. Sending 2Mbps point-to-point UDP/IP traffic (256 bytes)




Fig. 6. Receiving 2Mbps point-to-point UDP/IP traffic (256 bytes)

configurable) "RTS threshold", the sender simply senses the channel prior to sending the DATA message. The receiver also senses the channel before sending an ACK; ACK's take priority over other traffic due to their shorter sensing interval. As with broadcast traffic, this technique is vulnerable to "hidden terminal" collision, with the risk increasing as the message size increases. The linear formulation is:

$$E_{point-to-point-send} = m_{send} \times size + b_{send(<thresh)}$$

$$E_{point-to-point-recv} = m_{recv} \times size + b_{recv(<thresh)}$$

In this case, the b values are expected to have some intermediate value between broadcast and point-to-point traffic. The overall effectiveness of this technique, taking into account the increased likelihood of collision and retransmission is outside the scope of this work.

Discard Traffic

As a consequence of operating in ad hoc mode, a network interface overhears all traffic sent by nearby hosts. It is therefore important to consider not only the energy consumption of sending and receiving traffic, but also the energy consumed by an interface when it processes point-to-point traffic that it will discard after determining that it is not the intended destination. This case is interesting not only because it is the most amenable




Fig. 7. Discarding traffic (source and destination).




Fig. 8. Discarding traffic (source only).

to energy conserving strategies, but also because it represents the “common case” – most of the traffic is for somebody else.

The non-destination host’s location with respect to both source and destination determines what part of the IEEE 802.11 protocol it overhears. Figure 7 shows the case of a discarding host in range of both the source and destination. The case of a non-destination node in range of the sender only is shown in Figure 8. A suitable machine configuration was found by walking around with laptops running ‘ping’ tests.

In the equations, S and D represent the case of being within transmission range of the sender and destination, respectively).

$$E = m_{disc} \times size + b_{non-dest(S,D)}$$

$$E = m_{disc} \times size + b_{non-dest(S,\emptyset)}$$

The values of $b_{non-dest}$ reflect the cost of processing the MAC protocol messages in order to avoid contention with a transmission that is intended for another destination. For non-destination hosts in range of the sender, the sign of m_{disc} indicates what kind of energy conservation strategy is used by the network interface. If $m_{disc} > 0$, in particular, if $m_{disc} = m_{recv}$, then the non-destination host effectively receives the traffic before discarding it. If $m_{disc} = 0$, then the non-destination host maintains the network interface in the idle state during the data transmission. If $m_{disc} < 0$, then the interface must employ some energy-conserving strategy based on the presence of uninteresting data on the media. (Recall that all energy costs are defined relative to the idle power consumption of the interface.)

The oscilloscope trace in Figure 9 shows the case of a non-destination host in range of both the source and destination. While data is being transmitted, the Lucent interface uses slightly less power than it does in idle mode. While the source is transmitting data, the non-destination host can neither send nor receive⁵ point-to-point traffic because the source and destination have reserved the media via the RTS/CTS exchange. The non-destination host could, in principle, receive a broadcast message sent by a fourth host that is out of range of both the source and




Fig. 9. Non-destination (2Mbps) host in range of both sender and receiver discarding 2Mbps point-to-point UDP/IP traffic (256 bytes)

destination. But because the non-destination host is in range of both senders, the packet has a high probability of being unintelligible due to collision. The non-destination host therefore loses little or no traffic by entering this reduced energy mode. The amount of energy that can be saved using this technique is dependent on the amount of time that the data transmission takes, as well as the amount of time and energy needed to return to the idle state.

The network configuration for case of a non-destination host in range of destination only is analogous that of the non-destination host in range of the sender (Figure 8). This is expressed as:

$$E = m_{none} \times size + b_{non-dest(S,D)}$$

The coefficient m_{none} refers to the network interface’s behavior while data is being transmitted, even though it is out of range of the sender. In principle, m_{none} need not be 0, as the interface is aware of the data transmission via the CTS message.

Promiscuous Mode Operation

A non-destination host operating its network interface in promiscuous mode eavesdrops on all point-to-point data traffic that it overhears. The coefficients are therefore a combination of the receive and discard cases. The data is treated as in the case of point-to-point receive, but the control traffic is treated as in the case of point-to-point discard.

In range of both sender and destination:

$$E = m_{recv} \times size + b_{non-dest(S,D)}$$

In range of sender only:

$$E = m_{recv} \times size + b_{non-dest(S,\emptyset)}$$

A host that is only within wireless range of the destination cannot overhear the data that is being transmitted, even if it is operating in promiscuous mode. So this is the same as the discard case:

$$E = m_{none} \times size + b_{non-dest(S,D)}$$

⁵It can’t send the CTS/ACK.

Constant Values

The measurements used to determine the constants V_{in} and P_{idle} are presented in Table I.

Though not used in energy consumption calculations, the average current draw while receiving and transmitting data were measured in order to compare with the specification. The 11Mbps card shows good agreement. For the 2Mbps card, the observed current draw during data transmit and receive was lower than the nominal value. The reason for the discrepancy is not clear and the specification does not indicate any details about the measurement.⁶

TABLE I
LUCENT IEEE 802.11 WAVELAN PC CARD CHARACTERISTICS

2Mbps		
	measured	spec
Sleep Mode	14mA	9 mA
Idle Mode	178 mA	n/a
Receive Mode	204 mA	280 mA
Transmit Mode	280 mA	330 mA
11 Mbps		
Sleep Mode	10mA	10 mA
Idle Mode	156 mA	n/a
Receive Mode	190 mA	180 mA
Transmit Mode	284 mA	280 mA
Power Supply	4.74 V	5 V

Results

Table II shows the complete energy consumption results, specifying linear coefficients for each relationship described in Section V. The equivalent graphical representation in Figure 10 better shows the qualitative overview that we want to emphasize.

A number of interesting points:

1. Sending point-to-point (a) and broadcast (b) traffic have the same incremental cost, but point-to-point traffic has a higher fixed cost associated with the IEEE 802.11 control protocol. This is exactly as expected.
2. Receiving point-to-point (c) and broadcast (d) traffic show differ only in their fixed costs: Receiving point-to-point traffic has a high fixed cost, due to the cost of sending two control messages. Receiving broadcast traffic has the lowest fixed cost, representing the MAC header in the DATA message and physical overhead. This is also as expected.
3. Receiving broadcast traffic (d) and receiving traffic in promiscuous mode while in wireless range of the sender (e, g) were expected to show the same incremental cost. They were also expected to have slightly different fixed costs, as the hosts overhear a different portion of the control sequence depending on whether they are in wireless range of the destination, while broadcast traffic has no control sequence. The data show that

⁶As non-technical speculation, we note that the measured values for the older “Bronze” card are in fairly good agreement with those specified for the newer Lucent “Silver” IEEE 802.11 cards.

⁷Excluding only MAC and PLCP headers.

TABLE II
LUCENT IEEE 802.11 2 MBPS WAVELAN PC CARD
2.4 GHz DIRECT SEQUENCE SPREAD SPECTRUM
LINEAR MODEL POWER CONSUMPTION MEASUREMENTS

	$\mu W \cdot sec / byte^7$	$\mu W \cdot sec$
point-to-point send (a)	1.9 $\times size$	+ 454
broadcast send (b)	1.9 $\times size$	+ 266
point-to-point rcv (c)	0.50 $\times size$	+ 356
broadcast rcv (d)	0.50 $\times size$	+ 56
non-destination $n \in \mathcal{S}, \mathcal{D}$		
promiscuous rcv (e)	0.39 $\times size$	+ 140
discard (f)	-0.61 $\times size$	+ 70
non-destination $n \in \mathcal{S}, n \notin \mathcal{D}$		
promiscuous rcv (g)	0.54 $\times size$	+ 66
discard (h)	-0.58 $\times size$	+ 24
non-destination $n \notin \mathcal{S}, n \in \mathcal{D}$		
promiscuous “rcv” (i)	0.0 $\times size$	+ 63
discard (j)	0 $\times size$	+ 56
idle (ad hoc) (k)	843 mW	
idle (BSS)	66 mW	

the three relations form a cluster; however, while (d, g) have approximately equal slopes, promiscuous mode receive (e) is a significant outlier.

4. Non-destination hosts in range of the sender enter a reduced power consumption mode while data is being transmitted (f, h). Because this mode has lower power consumption than the idle mode, the incremental cost of ignoring data is negative, with the two slopes approximately equal, nearly as expected. The fixed costs differ as each host discards a different portion of the control sequence: Each control message that is discarded appears to cost about $25 \mu W \cdot sec$, though it is not clear that the accuracy of the

5. For a non-destination host in range of the destination, but not the sender (i), there is no corresponding energy conserving strategy. Neither this host nor a similarly located host operating its network interface in promiscuous mode (j) overhear data because they are not in range of the sender. Because this host cannot receive point-to-point traffic, it would be possible to use the same strategy as in (f, h). However, non-destination nodes in range of only the destination are more likely than those in range of the sender to be able to receive other broadcast traffic without collisions (i.e. from nodes in range of neither sender nor destination).

6. A non-destination host must process the control traffic, regardless of whether it goes on to discard or eavesdrop on the ensuing data traffic. Non-destination hosts in range of both sender and destination (e, f), in range of the sender only (g, h) and in range of the destination only (i, j) were expected to (pairwise) have the same fixed costs. The data show that the fixed cost was noticeably higher for a host operating in promiscuous mode.

The cause of this discrepancy is not clear. It is possible that the expectation was not correct: The energy management strategy used by non-destination hosts may also reduce fixed costs in some way that is not clear from the scope trace. It is also possible that operating the network interface in promiscuous mode




Fig. 10. Experimental results: 2 Mbps card

makes other operations more expensive: however, it is known *not* to have a significant effect on the idle mode power consumption.

7. The relative magnitude of the idle energy consumption is indicated by (k), which shows the energy consumed by an idle interface during the time it takes to transmit or receive *size* bytes at 2 Mbps, naively calculated as $3.37 \mu W \cdot sec/byte$.

A. Higher Transmission Rates

The experimental results for the 11Mbps card are shown in equation form in Table III.

The higher data transmission rate for the 11Mbps “Silver” card does not lead to a fivefold reduction in energy consumption, or even to a fivefold increase in available bandwidth.

1. In order to ensure that the 11Mbps cards can operate in the presence of slower cards, their RTS/CTS control messages must be properly received by the slower cards. This means that the fixed overhead, which is dominated by channel access overhead, does not decrease significantly. The slight overall decrease may be due to some change in the newer card or due to experimental uncertainty.

2. In order to ensure that 11 Mbps cards can interoperate with slower 2Mbps cards in the same network, broadcast messages must also be sent at the low data rate. Thus incremental cost of sending traffic does not decrease when sending broadcast traffic

(b). The slight increase is attributed to experimental variation. The incremental cost of receiving broadcast traffic decreases by a factor of two. The reason for this is unknown.

3. Point-to-point traffic between 11Mbps cards can be sent at the higher data rate. The incremental cost of sending and receiving (a, c) point-to-point data decreases by more than a factor of four. This is the expected behavior. At this point, the cost of sending data is so low that sending an 800 byte packet consumes only twice as much energy as sending an empty one.

4. The energy conservation strategy observed in 2Mbps cards is not observed in 11Mbps cards. Because the transmission time is so short, the expected benefit would be quite small. The incremental cost of discarding data (f,h) is about the same as the cost of both receiving (c) and promiscuous mode operation(e,g). However, as with the 2Mbps cards, the fixed overhead associated with promiscuous mode operation is higher. (Interactions between 2Mbps and 11Mbps cards were not investigated, nor was the response to extremely long (fragmented) data.)

B. Retransmissions

During one series of measurements – of an 11Mbps card receiving in promiscuous mode – data that could be interpreted as indicating retransmissions was recorded. The source of the potential interference was not identified and the situation could not be reproduced. The data is reproduced in Figure 11.

TABLE III
LUCENT IEEE 802.11 11 MBPS WAVELAN PC CARD
2.4 GHz DIRECT SEQUENCE SPREAD SPECTRUM
LINEAR MODEL POWER CONSUMPTION MEASUREMENTS

	$\mu W \cdot sec / byte^8$	$\mu W \cdot sec$
point-to-point send (a)	.48 $\times size$	+ 431
broadcast send (b)	2.1 $\times size$	+ 272
point-to-point rcv (c)	0.12 $\times size$	+ 316
broadcast rcv (d)	0.26 $\times size$	+ 50
promiscuous rcv (e)	non-destination $n \in \mathcal{S}, \mathcal{D}$	
discard (f)	0.14 $\times size$	+ 97
	.11 $\times size$	+ 66
promiscuous rcv (g)	non-destination $n \in \mathcal{S}, n \notin \mathcal{D}$	
discard (h)	0.10 $\times size$	+ 70
	0.11 $\times size$	+ 42
promiscuous "rcv" (i)	non-destination $n \notin \mathcal{S}, n \in \mathcal{D}$	
discard (j)	0.0 $\times size$	+ 32
	0 $\times size$	+ 38
idle (ad hoc) (k)	741 mW	
idle (BSS)	48 mW	




Fig. 11. Retransmissions?

C. Discussion

In general, data are in good agreement with what might be expected based on the IEEE 802.11 protocol definition. Certainly, they meet the goal of indicating general orders of magnitude and relative energy consumption costs of various portions of the IEEE 802.11 protocol operating in ad hoc mode.

There are a number of potential sources of experimental error. As can be seen from the scope traces, the signal is rather noisy. Likewise, the input voltage and idle current show similar high frequency noise in addition to variation associated with data traffic. There also seemed to be some longer term (daily) variation in the baseline energy consumption; sometimes this amounted to as much as 1-2mV or 10-20mW.

The linear model is clearly appropriate: the correlation coefficient is over .99 in nearly every instance. Approximately 50 - 90 packet measurements, each with an uncertainty of about 7%, contributed to the calculation of each linear equation. The energy measurements at each measured packet size clustered fairly

neatly, two standard deviations in each data point was generally less (and often much less) than $\pm 12\%$. The standard error of the calculated coefficients was generally less than 5%.

The experiments were carried out in an unshielded environment, though efforts were made to avoid traffic on the test channel. The traces show a certain amount of apparently random noise; using the averaging technique described, there was generally little variation between successive measurements. Results which seemed to indicate the occurrence of retransmission were only observed on one occasion. The effect seemed to be quite obvious, increasing our confidence that other data was relatively unaffected.

One interesting possibility is that the amount of processing (energy) required to receive a packet depends significantly on received signal strength and/or ambient interference. The relatively high incremental cost of receiving data is partly due to extensive signal processing performed at the receiver. If this requires significant computation, it may also require more energy. This effect may be difficult to quantify without more specialized equipment and a more controlled wireless environment.

VII. NEW PERSPECTIVES ON ENERGY-AWARE PROTOCOLS FOR AD HOC NETWORKS

The data obtained above would have only a limited computational value if they did not suggest new perspectives on developing protocols for the ad hoc environment.

Energy as a Distinct Metric

The data clearly show that energy consumption and bandwidth utilization are substantively different metrics. The former must take into account that the energy consumption of receiving and discarding traffic is both substantial and disproportional to the number of packets or bytes transmitted. This emphasizes the differences between broadcast and point-to-point traffic. In addition, the fixed overhead cost of sending or receiving a packet is relatively high, while the incremental cost of data is relatively low. This makes the mix of packet sizes associated with a protocol an important for energy consumption behavior.

Data from experiments of the type described in this paper have been [4] and should continue to be used as an adjunct to bandwidth-oriented simulation studies of protocols intended to operate in the ad hoc environment.

Speed and Energy Efficiency

The data rates provided by inexpensive, commercially available wireless interfaces cards has increased significantly over the past year or two. This has been combined with modest improvements in energy efficiency at the hardware level.

However, the relationship between transmit speed and overall energy consumption is complex. The faster transmit and receive times have only limited impact on the per-packet energy consumption. The incremental component of the cost is roughly linear in the transmit rate. However, the fixed component, which tends to dominate, is much less affected. It includes carrier sense, transfer to and from the card and mode switching, as well as the PLCP and access control messages, which

are transmitted at the slower rate ⁹.

In addition, the reduced transmission range associated with higher data rates increases the number of hops required to deliver traffic to its destination in a multi-hop routing environment. However, it also decreases the number of neighbors which hear each transmission. Simulation studies such as those carried out in [4] are needed to further investigate the complex tradeoffs involved. The problem seems to be similar, but not identical, to that of selecting optimal transmit power.

Energy Efficient Cluster-based Routing Protocols

The most striking feature in all this data is the extremely high idle power consumption of a wireless network interface operating in ad hoc mode. The issue does not seem to be adequately addressed by current work in protocol design for ad hoc networks.

Most ad hoc routing protocols are what is referred to in [5] as “uniform” protocols, in which the protocol operates in the same way at all nodes. There is also a smaller class of “partitioning” protocols. These protocols impose a structure on the network, partitioning it into clusters, dominated by “cluster-heads” and connected by “gateways”. These hosts take on a special role in managing routing information. However, connectivity changes may require expensive recomputations of cluster membership. Cluster maintenance overhead has therefore been seen as a serious disadvantage for these protocols.

The extremely high energy cost of an idle interface operating in ad hoc mode suggests that there may be significant advantage emulating the energy management strategies of BSS mode. Clustering provides an “infrastructure” in which a subset of nodes can buffer traffic for their neighbors, which could remain in a low power consumption state. Because this designated subset can be highly dynamic, there certainly will be additional complexity in cluster maintenance, which would ideally be integrated with the polling operation. Some penalty in latency and possible packet loss also seems unavoidable. Despite these issues, modification of protocols such as [9] or [14] to support such a energy-management strategy may prove worthwhile.

VIII. CONCLUSIONS

The results of this simple series of experiments show that the energy consumption of an IEEE 802.11 wireless interface has a complex range of behaviors that are relevant to the design of network layer protocols – energy consumption is not synonymous with bandwidth utilization. Energy-aware protocol design and evaluation must consider factors such as the relative proportions of broadcast and point-to-point traffic, packet size and reliance on promiscuous mode operation.

Measurements of the energy consumed sending, receiving and discarding packets of various sizes are presented as collection of linear equations as well as a visual form which highlights general conclusions: Energy consumption associated with receiving data is not negligible and fixed overhead costs are very high. This data should prove helpful to protocol developers who wish to include energy-aware protocol design and analysis in their work.

⁹Not only would the cards be incompatible, but the 2Mbps traffic would interfere with 11Mbps traffic.

Large improvements in data transmission rates have a fairly limited effect on overall per-packet energy consumption because the MAC protocol and broadcast traffic must use lower transmission rates. The complex relationships among transmission speed, range and effective node density and their effect on bandwidth and energy utilization in a multi-hop network are worth further study.

When operating in ad hoc mode, the idle power consumption is significant, as hosts must maintain their network interfaces in idle mode in order to cooperate in maintaining the ad hoc routing fabric. In particular, partitioning routing protocols, which dynamically maintain a cluster-based “infrastructure” may be well-suited to apply some variant of the energy management techniques used in a base station environment.

We would like to thank the anonymous reviewers and our colleagues at SICS, especially Anders Holst, for their helpful comments. Thank you also to Sycard Technology.

REFERENCES

- [1] Jae-Hawn Chang and Leandros Tassiulas. “Energy Conserving Routing in Wireless Ad-hoc Networks.” Proceedings of IEEE Infocom 2000. Tel Aviv, Israel, March 2000.
- [2] Jun-Cheng Chen, Krishna M. Sivalingam, Prathima Agrawal and Shalinee Kishore. “A Comparison of MAC Protocols for Wireless Local Networks Based on Battery Power Consumption.” Proceedings of IEEE Infocom 1998. San Francisco, CA, March, 1998.
- [3] Keith I. Farkas, Jason Flinn, Godmar Back, Dirk Grunwald, and Jennifer Anderson. “Quantifying the Energy Consumption of a Pocket Computer and a Java Virtual Machine” Proceedings of the International Conference on Measurement and Modeling of Computer Systems (ACM SIGMETRICS). Santa Clara, CA, June, 2000.
- [4] Laura Marie Feeney. “An Energy-consumption Model for Performance Analysis of Routing Protocols for Mobile Ad Hoc Networks.” To appear in *Journal of Mobile Networks and Applications*, 2001.
- [5] Laura Marie Feeney. “A Taxonomy for Routing Protocols in Mobile Ad Hoc Networks.” SICS Technical Report T99:07.
- [6] Laura Marie Feeney. “Investigating the Energy Consumption of an IEEE 802.11 Network Interface.” SICS Technical Report T99:11.
- [7] Paul Gauthier, Daishi Harada and Mark Stemm. “Reducing Power Consumption for the Next Generation of PDAs: It’s in the Network Interface.” Proceedings of MoMuC’96, September, 1996.
- [8] Wendi Rabiner Heinzelman, Joanna Kulik and Hari Balakrishnan. “Adaptive Protocols for Information Dissemination in Wireless Sensor Networks.” Proceedings of MobiCom’99, Seattle, WA, August, 1999.
- [9] Mingliang Jiang, Jinyan Li, Yong Chiang Tay. Cluster Based Routing Protocol Functional Specification. Internet Draft, draft-ietf-manet-cbrp-spec*.txt. Work-in-progress.
- [10] Robin Kravets and P. Krishnan. “Power Management Techniques for Mobile Communication.” Proceedings of MobiCom’98. Dallas, TX, October 1998.
- [11] P. Krishna, N. H. Vaidya, M. Chatterjee and D. K. Pradhan. “A Cluster-Based Approach for Routing in Dynamic Networks.” ACM SIGCOMM Computer Communications Review (CCR), 1997.
- [12] Lucent Technologies. “WaveLan IEEE 802.11 PC Card User’s Guide”. February, 1999. <http://www.wavelan.com>.
- [13] R. Ramanathan and R. Rosales-Hain. “Topology Control of Multihop Wireless Networks using Transmit Power Adjustment.” Proceedings of IEEE Infocom 2000. Tel Aviv, March 2000.
- [14] P. Sinha, R. Sivakumar, and V. Bharghavan, “CEDAR: a Core-Extraction Distributed Ad hoc Routing algorithm.” Proceedings of IEEE Infocom 1999, New York, NY. March 1999.
- [15] Sycard Technology. “PCCextend 140 CardBus Extender User’s Manual (Preliminary)”. July, 1996. <http://www.sycard.com>.
- [16] Transmeta. <http://www.transmeta.com/>.