When Underwater Acoustic Nodes Should Sleep With One Eye Open: Idle-time Power Management In Underwater Sensor Networks

Albert F. Harris III
University of Padova

Milica Stojanovic
MIT

Michele Zorzi
University of Padova

Categories and Subject Descriptors

C.2.2 [**Computer-Communication Networks**]: Network Protocols

General Terms

Computer network performance

Keywords

Underwater acoustic networking, energy efficiency, wakup/sleep modes

1. INTRODUCTION

The current interest in underwater sensor networks stems from the potential to use long term sensing devices to monitor the large mass of oceans on the planet (*e.g.*, underwater seismic event monitoring or underwater oil rig monitoring). To accomplish this, the sensor nodes must have the ability to self-configure the communication network and provide energy-efficient data transmission. To this end, researchers have begun devising MAC-layer protocols that minimize energy consumption for data transmission.

Acoustic modems typically present a number of modes of operation, similar to radio interfaces (*e.g.*, transmit, receive, sleep, etc.), each of which consumes different levels of energy. In radio communications, the cost of keeping the interfaces idle is high; therefore, a number of idle-time power management solutions have been devised (*e.g.*, GAF [1], STEM [2], TITAN [3]) to conserve energy during times of no communication. It is natural to attempt to use these same methods for energy conservation in underwater sensor networks. However, there are significant differences between acoustic modems and radios transceivers, making it doubtful whether previous conclusions will be valid for the underwater environment.

The relative costs of various interface modes are significantly different for acoustic devices than for radios. Typical radio interfaces [4] have similar costs for transmitting, receiving and idling. On the other hand, acoustic modems have very high transmission costs with respect to receive costs, and have very low idle costs. This implies that certain trade-offs worthwhile for radios may be too costly for acoustic modems. Furthermore, capabilities inherent in acoustic modems (*e.g.*, the possibility of an ultra-low power receive state) may cause solutions that were too expensive for radio to be justifiable in an underwater network.

The main contribution of this work is a preliminary evaluation of idle-time power management techniques for underwater sensor networks. Through an analysis of the energy consumption of various modes for acoustic modems, we show that for sensors that transmit data with a period on the order of minutes to a few hours, idle-time power management techniques that increase the needed transmission time may perform poorly. As an alternative, we investigate the use of a wakeup mode. Wakeup modes for radios are not a new idea, but they have not yet been widely adopted due to the fact that their implementation requires new hardware and this technology may not be mature enough. However, in this paper we argue for the use of wakeup modes in acoustic modems. To this end, we present an evaluation of four protocols via simulation, demonstrating that the use of an ultra-low power wakeup mode consistently results in the greatest energy savings.

The rest of this paper is organized as follows. Section 2 presents the characteristics of acoustic modems and presents their impact on idle-time protocols. Section 3 presents our evaluation of these protocols over different network traffic patterns for acoustic modems. Finally, Section 4 presents some conclusions and future directions.

2. ACOUSTIC MODEMS

Today's acoustic modem technology includes commercially available modems (*e.g.*, the Teledyne-Benthos modem [5]), as well as those developed for research purposes, such as the Woods Hole Oceanographic Institution's (WHOI) modem. We shall use this modem as an example to summarize acoustic transmission parameters relevant for the present study. A detailed description of the WHOI modem can be found in [6].

The WHOI acoustic modem has two basic modes of operation: low rate and high rate. Low rate transmission/detection is accomplished using FSK modulation and noncoherent detection, with a bit rate of 80 bits per second (bps). High rate transmission is accomplished using PSK modulation and coherent detection, with a variable bit rate between 2,500 and 5,000 bps.

The modem includes the main processor and the co-processor, which perform the signal processing functions needed at the physical layer and the MAC layer in the current implementation. The modem is coupled to the transducer, where electrical signals are converted into acoustical ones and vice-versa.

The main processor is used to generate the signals for transmission, and to receive the low rate signals. Detection of high rate signals requires adaptive equalization and multichannel combining,

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

WUWNet'06, September 25, 2006, Los Angeles, California, USA.

Copyright 2006 ACM 1-59593-484-7/06/0009 ...$5.00.

which are computationally intensive operations. These functions are implemented in the co-processor, which is engaged only when the modem is receiving high-rate signals.

The modem can be in one of the following states, each of which is characterized by different power consumption: **Transmit:** to transmit, the modem consumes between 10 and 50 watts, depending on the distance. For example, at 50 W, an acoustic signal power of 185 dB re μPa can be generated, which is sufficient for transmission over several kilometers in shallow water [6]. **Listen:** when in the listening state, the modem consumes 80 mW. In this state, the modem is waiting for a packet. A packet arrival is detected by receiving a packet preamble. The packet preamble also contains the information on the type of signal that is following, such as type of modulation, packet length, etc. **Receive, low rate:** to receive a data packet modulated using FSK (low rate) the modem consumes 80 mW. The processor performs noncoherent detection in this case, which requires no more power than needed for active listening. **Receive, high rate:** to receive a data packet modulated using PSK (high rate), the modem consumes 3 W. The co-processor must be engaged to perform coherent signal detection in this case, which requires more power than needed for noncoherent detection. **Sleep:** the modem is turned off in this state and is not capable of detecting signals.

Card	Transmit	Receive	Idle	Sleep
Cisco Aironet [4]	2240	1350	1350	75
Micro Modem [6]	10,000	3,000	80	≈ 0

Table 1: Power consumption (mW) for interface modes

It can easily be seen that the energy profiles of acoustic modems are quite different from those of typical radio transceivers. For the acoustic devices, the transmit power is an order of magnitude larger than that of receive mode. Furthermore, the idle mode is very cheap. Compare this to the typical energy consumption of a radio, where the transmit and receive powers are of the same magnitude, and the idle costs are very expensive (see Table 1).

An additional possibility is to use an ultra-low power wakeup state. Although some modems may have multiple low-power modes (*e.g.*, Teledyne-Benthos [5]), using such modes as a wakeup mechanism (in which a node in a low-power state continuously listens for a wakeup signal) is not currently done. However, it would be possible to do so should the overall power savings indicate the benefits of such an approach. As an example, Heidemann *et al.* [7] have begun providing such a mode in a new modem, which is capable of detecting a signal by monitoring the acoustic energy in the channel. These modes have not yet been used for the purpose of idle-time power management as we propose here however, nor have the benefits of doing so been quantified for acoustic modems. The circuit in [7] uses a dual gate FET configured as a cascade amplifier, with a passive filter and detector, and consumes $500\,\mu W$. When in this state, the modem would wait for a wakeup signal. Upon receiving the signal, the modem would transition into listening state and from there into active reception.

Switching from one state to another happens almost instantaneously, except for several hundred milliseconds that are needed to power up the co-processor. No extra power is required to switch from one state to another.

The large difference in the power needed to transmit an acoustic signal and that needed to receive and process it, motivates the search for a suitable MAC protocol for use in an underwater sensor network. Two of the main performance metrics for MAC protocol evaluation are throughput efficiency and energy efficiency. While the throughput efficiency remains fundamentally limited by the long propagation delay of acoustic signals [8, 9], significant savings in energy consumption can be obtained through minimizing the amount of time the modem spends in transmit mode. Minimizing the energy consumption is especially important in underwater networks of fixed nodes, which are battery-powered and intended for long-term deployment.

Although the applications of underwater sensor networks are still evolving, one can envision at least two types of applications: event-driven and periodic sensing. The two types of applications imply different traffic patterns. In this work, we focus on a network of sensors whose task is to constantly sense their environment and report their findings to an end node. The rate at which the information is generated (*i.e.*, the number of packets per second per node and the node density) determines the level of network activity that must be supported. In this work, we analyze and compare four different protocols for varying traffic generation rates.

2.1 Sleep Cycling

It has been suggested [7] that underwater sensor networks should have supernodes every few tens of nodes to help minimize the time for data collection, depending on the application. Networks of mobile unmanned vehicles will likely be even more sparse, due to the high cost of building and deploying them. This poses an immediate difference with radio networks. Each node in an underwater sensor network is likely to be vital to the connectivity of the network. Therefore, any method that attempted to keep a backbone awake at all times would likely have all of the nodes awake 100% of the time. Furthermore, any sort of randomized wakeup sequences would also perform poorly due to this expected low node density.

On the other hand, methods that build paths on-demand also are not ideal. First, most of these schemes increase the delay until a node can receive data. The effects of this sort of delay increase are magnified in an event driven network, where timely delivery of packets could be critical. Second, many of these schemes require a sender to transmit a wakeup beacon in such a way that it is guaranteed to be received, often by repeated transmission. But for acoustic modems, transmission is much more expensive than any other mode, causing such beaconing to potentially outweigh the savings gained by being in sleep mode.

Essentially, any on-demand scheme must have a way to wake up a sleeping node. Most of these schemes use some type of low duty cycle wakeup for nodes to listen for incoming transmissions [2, 3]. Senders are required to transmit a beacon, or request to transmit, in such a way that the intended receiver is guaranteed to hear it.

Consider a sleep cycle where T_{rx} is the time that a receiver is listening. Then it is clear that only if the beacon falls within T_{rx} will the node be successfully awakened. For a given interval T, $T_{sleep} = T - T_{rx}$. Let the beacon be of length B and the inter-beacon time be B_l (the receiver must respond in this time). Schurgers *et al.* [2], show that the average time a sender will spend sending beacons (T_b) is as follows: $T_b = \frac{T+(B+B_l)}{2}$. This demonstrates a basic trade-off between the amount of time spent sleeping and the amount of time spent sending beacons. For acoustic modems, where the transmit energy consumption is so large, these beaconing periods can consume a large amount of energy.

Consider the case where $T_{rx} = 225ms$ and $B + B_l = 150ms$. For the node to sleep for 75% of the idle time, the average time it will be sending beacons is nearly 300 ms [2]. These numbers are reasonable for radio networks but would be larger for acoustic modems due to the increased latencies, having the effect of further increasing the energy consumption. Even at the lowest transmit power of 10 W, this 300 ms transmission for the sender and a 75 ms transmission for the receiver translates to 3,750 mJ consumed to wake up the node. This is nearly one minute of standard idle time; therefore, if the generated traffic is about a packet a minute or more, there is no benefit in adopting a sleep cycle of this kind. Now, consider the possibility of having the ultra-low power wakeup mode consuming only $500\ \mu W$. The energy spent beaconing then translates to over 2 hours of wakeup mode time, making the wakeup protocol even more advantageous, except for very low traffic scenarios. In our numerical results, we will use a CSMA-based MAC protocol. A detailed comparison among different MAC schemes (including scheduled TDMA-based MAC) is left for future research, as in this paper we focus on evaluating the potential for energy savings via sleep modes or wakeup modes rather than on the optimization of the MAC protocol used by the nodes when they are awake.

2.2 Acoustic Wakeup

The ability of acoustic modems to implement an ultra-low power wakeup state yields another option. In the case of radio, the extra hardware and difficulties in implementation may outweigh the benefits; however, for certain traffic patterns, we expect such a mode would yield significant savings over sleep cycling methods. Essentially, the amount of energy saved by transitioning into a low power sleep mode must outweigh any energy expended to wake up intended receivers for asynchronous sleep cycling solutions to be efficient. Because transmit power is so high for acoustic modems and idle energy is so low, this sleep time must be significantly longer than for radio sensor networks.

Additionally, implementing wakeup modes in acoustic modems is considerably easier. First, no extra transducer is needed, reducing the cost of implementation. Recall from early in this section that a $500\ \mu W$ wakeup mode can be realized using very simple decoding [7]. In principle, it is possible to design a signal that requires only very simple processing. This type of signal is likely to rely on a set of tones, or a chirp, that are amenable to low-complexity processing.

In the next section we compare the effects of network traffic patterns on the energy efficiency of various sleep mechanisms. The results demonstrate that it is worthwhile to implement wakeup modes in acoustic modems given the significant energy savings achievable over traditional sleep cycling solutions.

3. ENERGY ANALYSIS

The goal of the following evaluations is to determine when a wakeup state is preferable to a sleep cycling solution for underwater sensor networks. To this end we compare four protocols: **Standard Idle** simply stays in idle state and never transitions to a sleep or wakeup mode. **Optimal Sleep** transitions immediately into a sleep mode and only wakes up during active transmission and reception. **STEM [2]** uses a sleep schedule, as described in Section 2.1 for receivers to transition in and out of sleep mode. If a wakeup signal is received, the receiver sends a "ready to receive" message to the transmitter and transitions into the active listening state. **Wakeup Mode** transitions into an ultra-low power wakeup mode after transmission and reception.

There are a number of ways to evaluate the impact of protocols on energy consumption in a sensor network. One method is to evaluate the total energy consumption in the network for various traffic patterns. Another method is to evaluate the time to first node death (or more generally the time until a given percentage of nodes die), which corresponds to evaluating the maximum energy consumption across nodes. The time to network partition is a metric that is closely related to the latter, insofar as the first node deaths often correspond to network partition in sparse networks. We choose to look at both total energy consumption and time to first node death in this study, with time to network partition being left for a future expanded study.

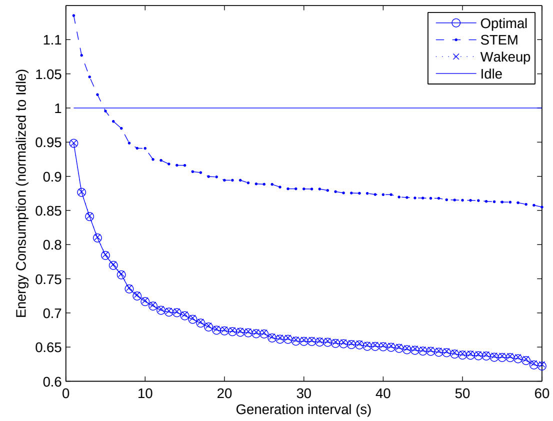

Figure 1: Total energy consumption of the network vs. traffic generation interval.

3.1 Simulation Setup

We used the ns2 simulator [10] to run our experiments. To account for energy consumption, ns2 is augmented with an energy model of the four protocols in various states using the values in Section 2 with a 10 W transmit power. The network covers an 1000 m by 1000 m area, in which 25 nodes are deployed randomly. We further modified the ns2 physical layer and propagation model to approximate the properties of the WHOI acoustic modem. A CSMA MAC layer is used and routing is done via directed diffusion [11]. For our evaluations, we use the average of 20 runs for each set of parameters tested. The resulting 95% confidence intervals are within $\pm 2\%$ of the values shown.

3.2 Evaluation

In this section, we evaluate the performance, in terms of energy consumption, of the four protocols discussed above in two different situations: under different traffic generation rates with a fixed wakeup mode cost of 500μ W, and as the cost of the wakeup mode increases.

As the interval between events in the network increases, the amount of possible sleep time increases. Therefore, idle-time power management solutions should save larger amounts of energy for longer traffic generation intervals. Figure 1 shows the energy consumption of the entire network for each of the four protocols as the interval between sensing events ranges from one second to one minute per node. Each value is normalized to the energy consumption of the entire network for the standard idle protocol. As can be seen, the wakeup mode protocol performs almost optimally. This is because the wakeup mode consumes almost no energy and does not require any additional transmission. STEM, however, due to the probability that a wakeup signal will be transmitted for some portion of the sleep interval, saves significantly less energy. Similar curves for times up to 4 hour intervals were roughly the same (*e.g.*, for a four hour interval, STEM: 0.76, Wakeup: 0.55, Optimal: 0.54), with STEM always consuming more energy due to increased transmission times. It is worth pointing out that this represents a worst-case for idle management solutions since in such a sparse network, virtually all nodes are needed for forwarding traffic.

The primary reason why STEM performs so poorly is that the

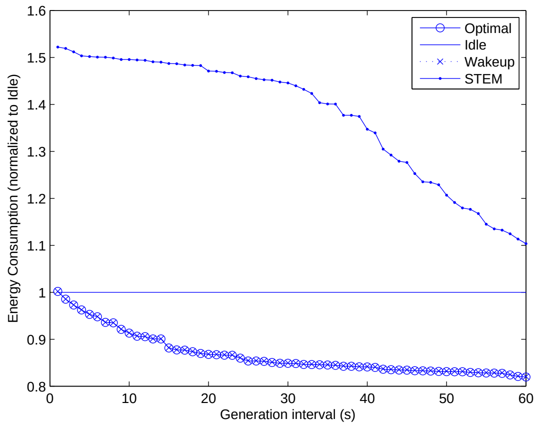

Figure 2: Highest energy consumption of a node vs. traffic generation interval.

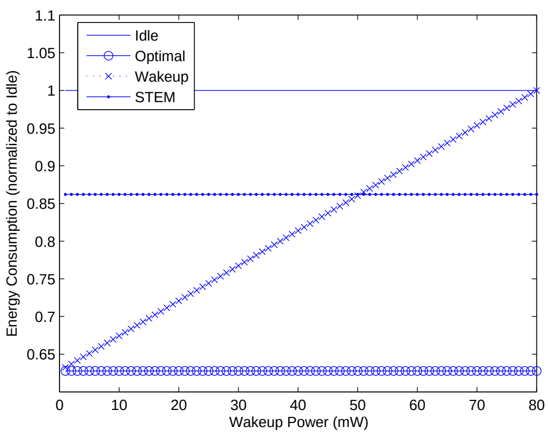

Figure 3: Total energy consumption of the network vs. wakeup mode cost.

transmit mode energy consumption of the acoustic modem is so high (in this case 10 W) that sending the wakeup beacon is very costly. Therefore, nodes that send the most traffic have much greater costs than the rest of the nodes. The greatest amount of energy consumed by a node is depicted in Figure 2. Increasing a single node's energy consumption is another definite drawback of any sleep cycling solution that increases the transmission time needed to send data. As can be seen in this figure, certain nodes have their energy expenditure increased dramatically over the average network energy consumption. This will lead to rapid node failure. If the underwater sensor networks are sparse, then this will rapidly result in network segmentation. Using a wakeup mode again keeps the energy consumption very close to optimal.

The main reason why the wakeup mode protocol performs so near optimal for these situations is the extremely low power used. A fair question to explore is: How low does this power have to be? To answer this we again look at the same scenario, but this time fix the sensor event frequency at once per minute per node and vary the power of the wakeup mode between 1 mW and 80 mW (the cost of idle mode). Figure 3 depicts the total energy consumption of the network. For this traffic rate, the wakeup mode protocol outperforms STEM for powers lower than about 50 mW. Recall the 500μ W figure used early, even if this number were to be increased by a factor of 10, there would still be very significant gains. As the time between sensor events increases, the cross-over value decreases; however, for events happening more often than every few hours, the wakeup mode still has the potential to outperform STEM. Furthermore, to accurately implement a solution like STEM, information about traffic generation rates is used to optimize the sleep cycle. This information may not be available in highly dynamic environments. The use of a wakeup mode avoids the need for such information, making the proposed solution more flexible and robust.

4. CONCLUSIONS

This paper has examined how the differences between acoustic modems and radios affect the design of idle-time power management schemes. Because idle-time power management schemes that use asynchronous sleep cycling trade off increased transmission time for increased sleep time, their performance when faced with the extremely high transmit power costs in acoustic modems may be very poor.

A possibility to implement an ultra-low power wakeup mode in acoustic modems would offer an alternative to idle-time sleep cycling. We show through analysis and simulation that for underwater sensor networks where the expected traffic generation is less than one packet per node per few hours, the wakeup mode will save energy over sleep cycling both in terms of total network energy consumed and in terms of the greatest energy consumption of a single node.

5. REFERENCES

[1] Y. Xu, J. Heidemann, and D. Estrin, "Geography-informed energy conservation for ad hoc routing," in *Mobicom*, 2001.

[2] C. Schurgers, V. Tsiatsis, S. Ganeriwal, and M. Srivastava, "Optimizing sensor networks in the energy-latency-density design space," *IEEE Transactions on Mobile Computing*, vol. 1, no. 1, pp. 70–80, January-March 2002.

[3] C. Sengul and R. Kravets, "Conserving energy with on-demand topology management," in *MASS*, November 2005.

[4] Cisco, "Cisco Aironet 350 client data sheet," http://www.cisco.com/.

[5] Teledyne, "Teledyne-benthos modem," http://www.rdinstruments.com/nemo.

[6] L. Freitag, M. Grund, S. Singh, J. Partan, P. Koski, and K. Ball, "The WHOI micro-modem: An acoustic communications and navigation system for multiple platforms," http://www.whoi.edu, 2005.

[7] J. Heidemann, W. Ye, J. Willis, A. Syed, and Y. Li, "Research challenges and applications for underwater sensor networking," in *IEEE WCNC*, 2006.

[8] M. Molins and M. Stojanovic, "Slotted FAMA: a MAC protocol for underwater acoustic networks," in *Proc. IEEE Oceans Conference*, 2006.

[9] M. Stojanovic, "Optimization of a data link protocol for underater acoustic networks," in *Proc. IEEE Oceans Conference*, 2005.

[10] ns2 Network Simulator, http://www.isi.edu/nsnam/ns/.

[11] C. Intanagonwiwat, R. Govindan, and D. Estrin, "Directed diffusion: A scalable and robust communication paradigm for sensor networks," in *MobiCom*, 2000.