Computer Science Technical Reports   Computer Science

Spring 5-1-2006

X-MAC: A Short Preamble MAC Protocol for Duty-Cycled Wireless Sensor Networks ; CU-CS-1008-06

Michael Buettner
University of Colorado Boulder

Gary Yee
University of Colorado Boulder

Eric Anderson
University of Colorado Boulder

Richard Han
University of Colorado Boulder

Follow this and additional works at: http://scholar.colorado.edu/csci_techreports

X-MAC: A Short Preamble MAC Protocol For Duty-Cycled Wireless Sensor Networks

Michael Buettner, Gary Yee, Eric Anderson, Richard Han

Department of Computer Science
University of Colorado at Boulder

May 2006

X-MAC: A Short Preamble MAC Protocol For Duty-Cycled Wireless Sensor Networks[*]

Michael Buettner, Gary Yee, Eric Anderson, Richard Han
Department of Computer Science, University of Colorado, Boulder. Boulder, CO [USA]
{michael.buettner, gary.yee, eric.anderson, richard.han}@colorado.edu

Abstract

In this paper we present X-MAC, a low power MAC protocol for wireless sensor networks (WSNs). Standard MAC protocols developed for duty-cycled WSNs such as BMAC, which is the default MAC protocol for TinyOS, employ an extended preamble and preamble sampling. While this "low power listening" approach is simple, asynchronous, and energy-efficient, the long preamble introduces excess latency at each hop, is suboptimal in terms of energy consumption, and suffers from excess energy consumption at non-target receivers. X-MAC proposes solutions to each of these problems by employing a shortened preamble approach that retains the advantages of low power listening, namely low power communication, simplicity and a decoupling of transmitter and receiver sleep schedules. We demonstrate through implementation and evaluation in a wireless sensor testbed that X-MAC's shortened preamble approach significantly reduces energy usage at both the transmitter and receiver, reduces per-hop latency, and offers additional advantages such as flexible adaptation to both bursty and periodic sensor data sources.

1   Introduction and Motivation

Energy efficiency is a fundamental theme pervading the design of communication protocols developed for wireless sensor networks (WSNs), including routing and MAC layer protocols. One of the primary mechanisms for achieving low energy operation in energy-constrained WSNs is duty cycling. In this approach, each sensor node periodically cycles between an awake state and a sleep state. Key parameters that characterize the duty cycle include sleep time, wake time, and the energy consumed during the awake state and the sleep state. The period of a duty cycle is equivalent to its sleep time plus awake time. Given duty cycling sensor nodes, the challenges faced by designers of communication protocols are how to achieve high throughput, low delay, and energy efficiency as nodes are waking and sleeping in the network. This paper focuses on the design of X-MAC, an adaptive energy-efficient MAC layer protocol for duty-cycled WSNs.

Standard MAC protocols developed for duty-cycled WSNs can be roughly categorized into synchronized and asynchronous approaches, along with hybrid combinations.

These approaches are motivated by the desire to reduce idle listening, which is the time that the node is awake listening to the medium even though no packets are being transmitted to that node. Idle listening has been found in 802.11 protocols to consume substantial energy [8, 15], and therefore must be avoided in WSNs. Synchronized protocols, such as S-MAC [17] and T-MAC [16], negotiate a schedule that specifies when nodes are awake and asleep within a frame. Specifying the time when nodes must be awake in order to communicate reduces the time and energy wasted in idle listening. Asynchronous protocols such as B-MAC [14], and WISEMAC [7], rely on *low power listening*, also called preamble sampling, to link together a sender with data to a receiver who is duty cycling. Idle listening is reduced in asynchronous protocols by shifting the burden of synchronization to the sender. When a sender has data, the sender transmits a preamble that is at least as long as the sleep period of the receiver. The receiver will wake up, detect the preamble, and stay awake to receive the data. This allows low power communication without the need of explicit synchronization between the nodes. The receiver only wakes for a short time to sample the medium, thereby limiting idle listening. Hybrid protocols also exist that combine a synchronized protocol like T-MAC with asynchronous low power listening [9].

A key advantage of asynchronous low power listening protocols is that the sender and receiver can be completely decoupled in their duty cycles. The simplicity of this design removes the need for, and the overhead introduced by, synchronized wake/sleep schedules. Studies of lower power listening have demonstrated its energy-saving capabilities [14, 9].

While the low power listening approach is simple, asynchronous, and energy-efficient, the long preamble in low power listening exhibits several disadvantages: it is suboptimal in terms of energy consumption at both the sender and receiver; it is subject to overhearing that causes excess energy consumption at non-target receivers; and it introduces excess latency at each hop. First, the receiver typically has to wait the full period until the preamble is finished before the data/ack exchange can begin, even if the receiver has woken up at the start of the preamble. This wastes energy at both the receiver and transmitter. Second, the low power listening approach suffers from the *overhearing problem*, where receivers who are not the target of the sender also wake up

during the long preamble and have to stay awake until the end of the preamble to find out if the packet is destined for them. This wastes energy at all non-target receivers within transmission range of the sender. Third, because the target receiver has to wait for the full preamble before receiving the data packet, the per-hop latency is lower bounded by the preamble length. Over a multi-hop path, this latency can accumulate to become quite substantial.

This paper proposes a new approach to low power listening called X-MAC, which employs a *short preamble* to further reduce energy consumption and to reduce latency. The first idea is to embed address information of the target in the preamble so that non-target receivers can quickly go back to sleep. This addresses the overhearing problem. The second idea is to use a *strobed preamble* to allow the target receiver to interrupt the long preamble as soon as it wakes up and determines that it is the target receiver. This short strobed preamble approach reduces the time and energy wasted waiting for the entire preamble to complete. We demonstrate through implementation in a wireless sensor testbed that X-MAC results in significant savings in terms of both energy and per-hop latency. Finally, X-MAC includes an automated algorithm for adapting the duty cycle of the nodes to best accomodate the traffic load in the network. We demonstrate the additional savings in energy and latency achieved by this adaptation.

This paper makes the following contributions:

- X-MAC introduces a series of short preamble packets each embed target address information, thereby avoiding the overhearing problem of low power listening, and saving energy on non-target receivers.

- X-MAC inserts pauses into the series of short preamble packets, creating a strobed preamble, which enables the targeted receiver to shorten the strobed preamble via an early acknowledgement, thereby achieving additional energy savings at both the sender and receiver, as well as a reduction in per-hop latency.

- We describe an adaptive algorithm for automatically adjusting the duty cycle of receivers to the offered traffic load, which further reduces per-hop latency.

- Experimental evaluation validates the performance gains and energy savings of the X-MAC protocol in comparison to a traditional asynchronous duty cycle techniques.

In the following, Section 2 describes related work. Section 3 describes the basic X-MAC protocol design. Section 4 presents an optimal algorithm for adapting the receiver's duty cycle, and then presents a practical approximation. Section 5 describes the experimental implementation and evaluation on a testbed of motes. Sections 6 and 7 provide a discussion of future work and our conclusions.

2 Related Work

There are a number of approaches to duty-cycling MAC protocols seen in the literature. These approaches can be broadly divided into two categories: techniques that use some method of synchronization to assure that the wake periods of the nodes are concurrent; and those that have no synchronization requirements and instead depend on an extended preamble and low power listening.

S-MAC [17] is a low power RTC-CTS based MAC protocol that makes use of loose synchronization between nodes to allow for duty cycling in sensor networks. The protocol uses three techniques to achieve low power duty cycling: periodic sleep, virtual clustering, and adaptive listening. The nodes in the network periodically wake up, receive and transmit data, and return to sleep. At the beginning of the awake period, a node exchanges synchronization and schedule information with its neighbors to assure that the node and its neighbors wake up concurrently. This schedule is only adhered to locally, resulting in a virtual cluster, which mitigates the need for system wide sychronization. Nodes that lie on the border of two virtual clusters adhere to the schedules of both clusters, which maintains connectivity across the network. After the synchronization information is exchanged, the nodes transmit packets using RTS-CTS until the end of the awake period and the nodes then enter sleep mode. In [18], the authors introduce adaptive listening to reduce latency. With this, when a node hears an RTS or CTS from its neighbor, it will wake up briefly at the end of the transmission. If the node is the next hop on the data path, waking up at the end of the transmission will reduce latency as the packet can be forwarded immediately without having to wait until the next scheduled awake period.

T-MAC [16] improves on the design of S-MAC by shortening the awake period if the channel is idle. In S-MAC, the nodes will remain awake through the entire awake period even if they are neither sending nor receiving data. T-MAC improves S-MAC by listening to the channel for only a short time after the synchronization phase, and if no data is received during this window, the node returns to sleep mode. If data is received, the node remains awake until no further data is received or the awake period ends. The authors show that, for variable workloads, T-MAC uses one fifth of the energy used by S-MAC. While this adaptive duty cycling reduces energy usage for variable workloads, these gains come at the cost of reduced throughput and increased latency.

A comparison of duty cycling MAC protocols for WSNs is performed in [9]. Specifically, S-MAC and T-MAC are compared to standard CSMA/CA. S-MAC and T-MAC are also modified to use low power listening during the awake period, which further decreases the energy consumption of the protocols. While they show that T-MAC in combination with low power listening provides very low power communication, the protocol still suffers the same drawbacks as T-MAC, namely high latency and overhead associated with synchronization.

B-MAC [14], developed at the University of California at Berkeley, is a CSMA-based technique that utilizes low power listening and an extended preamble to achieve low power communication. Nodes have an awake and a sleep period, and each node can have an independent schedule. If a node wishes to transmit, it precedes the data packet with a preamble that is slightly longer than the sleep period of the receiver. During the awake period, a node samples the medium and if a preamble is detected it remains awake to

receive the data. With the extended preamble, a sender is assured that at some point during the preamble the receiver will wake up, detect the preamble, and remain awake in order to receive the data. B-MAC also provides an interface by which the application can adjust the sleep schedule to adapt to changing traffic loads. The method of adaptation is left to the application developer. The authors show that B-MAC surpasses existing protocols in terms of throughput, latency, and for most cases energy consumption. While B-MAC performs quite well, it suffers from the overhearing problem, and the long preamble dominates the energy usage.

WiseMAC [7], which is based on Aloha, also uses preamble sampling to achieve low power communications in infrastructure sensor networks. WiseMAC uses a similar technique to B-MAC, but the sender learns the schedules of the receiver awake periods, and schedules its transmission so as to reduce the length of the extended preamble. To achieve this, the receiver puts the time of its next awake period in the data acknowledgement frame. The next time the transmitter wants to send to that receiver it can begin the preamble only a short time before the receiver will awaken, taking into account possible clock skew. This reduces the energy expended when sending the preamble. In addition, for low traffic loads where the preamble is longer than the data frame, WiseMAC repeats the data frame in place of the extended preamble. Receivers process this data frame and if the node is not the intended recipient it returns to sleep. If the node is the recipient, it remains awake until the end of the transmission and sends an acknowledgement. While WiseMAC solves many of the problems associated with low power communications, it does not provide a mechanism by which nodes can adapt to changing traffic patterns.

In addition, low power listening has been implemented by a number of commercial radios, for example the Chipcon CC2500 [1] and the MaxStream XBee radios [3]. The XBee radio modules allow the user to set the sleep period of the radio and to set the length of the preamble that precedes the data packet. The user must be sure to set the sleep period to a duration shorter than the preamble length in order to be assured that the radio was awakened by the preamble. The Chipcon CC2500 uses a similar mechanism, but it has the added benefit of using a low power radio circuit that listens for the preamble. If in Wake-On-Radio mode, a low power radio circuit is used to intermittently sample the channel for a preamble. If the preamble is detected, the main radio circuit is woken up and the radio receives the data packet.

A variety of techniques have employed a Wake-On-Radio approach [15, 12] for energy-efficient communication. These approaches employ a second low power radio as a trigger to wake up the primary radio. These WOR approaches require special hardware assistance.

The 802.11 MAC protocol implements a power save mode that allows a base station access point to synchronize receivers to a wake/sleep schedule [13]. The base station transmits periodic beacons once every 100 ms to synchronize the clients. The clients wake up at the start of a beacon interval and checks the beacon to see if there is queued data waiting for it. If so, it stays awake and otherwise goes back to low power sleep mode. This is designed for infrastructure mode wireless LANs.

3 X-MAC Protocol Design

The design goals of the X-MAC protocol for duty-cycled WSNs are:

- energy-efficiency
- simple, low-overhead, distributed implementation
- low latency for data
- high throughput for data
- adaptivity to offered data load
- applicability across all types of packetizing and bit stream digital radios

For many applications, asynchronous duty cycling techniques are preferable to synchronized techniques in terms of energy consumption, latency, and throughput. In part, this is because they do not incur overhead due to synchronization. In addition, asynchronous techniques do not have to share schedule information and only stay awake long enough to sample the medium unless, of course, they are receiving or transmitting data. Hence, the awake period can be significantly shorter than that of synchronized methods. With a shorter awake period, asynchronous protocols can wake up more often while still maintaining a low duty cycle. Consequently, they experience reduced latency and higher throughput. However, as acceptable latency increases, the extended preamble begins to dominate energy consumption for asynchronous techniques. In general, for applications with very loose latency requirements, synchronized approaches may be more appropriate. In [14], the authors show that for a 10 hop network B-MAC outperforms S-MAC with respect to energy for latencies under 6 seconds.

For these reasons, X-MAC builds upon the foundation provided by asynchronous duty-cycled MAC protocols. While asynchronous techniques perform quite well, there are a number of problems which, if mitigated, would allow for even more efficient communication. X-MAC's design is motivated by the goal of mitigating the following four problems of low power listening: overhearing, excessive preambles, packetizing radios, and lack of automated adaptation to varying traffic loads.

3.1 Asynchronous Duty Cycling

A visual representation of asynchronous low power listening (LPL) duty cycling is summarized in the top section of Figure 1. When a node has data to send, it first transmits an extended preamble, and then sends the data packet. All other nodes maintain their own unsynchronized sleep schedules. When the receiver awakens, it samples the medium. If a preamble is detected, the receiver remains awake for the remainder of the long preamble, then determines if it is the target. After receiving the full preamble, if the receiver is not the target, then it goes back to sleep.

3.2 Embedding the Target ID in the Preamble to Avoid Overhearing

A key limitation of LPL is that non-target receivers who wake and sample the medium while a preamble is being sent
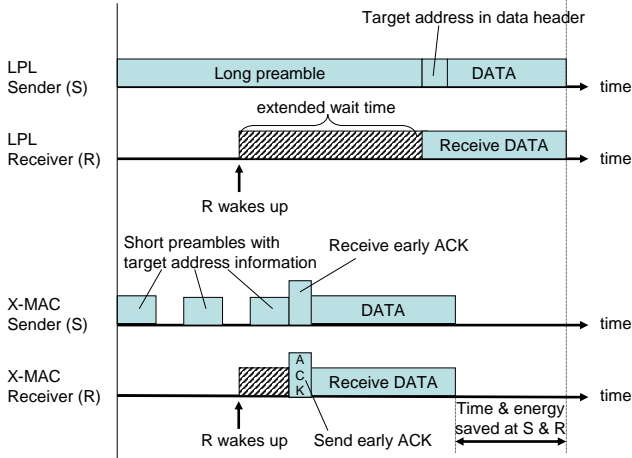
Figure 1. Comparison of the timelines between LPL's extended preamble and X-MAC's short preamble approach.

must wait until the end of the extended preamble before finding out that they are not the target and should go back to sleep. This is termed as the overhearing problem, and accounts for much of the inefficiency and wasted energy in current asynchronous techniques. This means that for every transmission, the energy expended is proportional to the number of receivers in range. Hence, the energy usage is dependent on density as well as traffic load. This problem is exacerbated by the fact that sensor networks are often deployed with high node densities in order to provide sensing at a fine granularity.

In X-MAC, we ameliorate the overhearing problem by dividing the one long preamble into a series of short preamble packets, each containing the ID of the target node, as indicated in Figure 1. The stream of short preamble packets effectively constitutes a a single long preamble. When a node wakes up and receives a short preamble packet, it looks at the target node ID that is included in the packet. If the node is not the intended recipient, the node returns to sleep immediately and continues its duty cycling as if the medium had been idle. If the node is the intended recipient, it remains awake for the subsequent data packet. As seen in the figure, a node can quickly return to sleep, thus avoiding the overhearing problem.

With this technique, the energy expenditure is independent of network density. The approach of a series of short preamble packets scales well with increasing density, i.e. as the number of senders increases in a neighborhood, energy expenditure remains largely flat. In comparison, as the number of senders increase in each neighborhood of a WSN practicing LPL, the entire WSN stays awake for increasing amounts of time.

Another advantage of this approach is that it can be employed on all types of radios. Any packetizing radio, such as the CC2420 characteristic of MICAZ and TelosB motes, the CC2500, and/or the XBEE, will be capable of sending a series of short packets containing the target ID. As we will

see later, such universal support across packetizing radios is not true of the traditional extended preamble LPL. In addition, the short preamble packets can be supported across all radios with bit streaming interfaces, e.g. the CC1000 that is found in the Mica2 mote.

3.3 Short Strobed Preamble to Reduce Excessive Preamble

Using an extended preamble and preamble sampling allows for low power communications, yet even greater energy savings are possible if the total time spent transmitting preambles is reduced. In traditional asynchronous techniques, the sender sends the entire preamble even though, on average, the receiver has woken up half way through the preamble. The entire preamble needs to be sent before every data transmission because there is no way for the sender to know that the receiver has woken up. This is one case where more time is spent sending the preamble than is necessary, as illustrated by the extended wait time in Figure 1. Another case occurs when there are a number of transmitters waiting to send to a particular receiver. After the first sender begins transmitting preamble packets, subsequent transmitters will stay awake and wait until the channel is clear. They will then begin sending their preamble, and this occurs for every subsequent sender. Consequently, each sender transmits the entire preamble when in fact the receiver was woken up by the first transmitter in the series.

In the development of X-MAC, we provide solutions for both of these cases. Instead of sending a constant stream of preamble packets, as would most closely approximate traditional LPL, we insert small pauses between each packet in the series of short preamble packets, during which time the transmitting node pauses to listen to the medium. These gaps enable the receiver to send an *early acknowledgement* packet back to the sender by transmitting the acknowledgement during the short pause between preamble packets. When a sender receives an acknowledgement from the intended receiver, it stops sending preambles and sends the data packet. This allows the receiver to cut short the excessive preamble, which reduces per-hop latency and energy spent unnecessarily waiting and transmitting, as can be seen in Figure 1. Since the sender quickly alternates between a short preamble packet and a short wait time, we term this approach a *strobed preamble*.

In addition to shortening the preamble by use of the acknowledgement, X-MAC also addresses the problem of multiple transmitters sending the entire preamble even though the receiver is already awake. In X-MAC, when a transmitter is attempting to send but detects a preamble and is waiting for a clear channel, the node listens to the channel and if it hears an acknowledgement frame from the node that it wishes to send to, the transmitter will backoff a random amount and then send its data without a preamble. The randomized backoff is necessary because there may be more than one transmitter waiting to send, and the random backoff will mitigate collisions between multiple transmitters. Also, the backoff is long enough to allow the initial transmitter to complete its data transmission. To enable this technique, after the receiver receives a data packet it will remain awake

for a short period of time in case there are additional transmitters waiting to send. The period that a receiver remains awake after receiving a data packet is equal to the maximum duration of the senders backoff period, to assure that the receiver remains awake long enough to receive any additional transmitters data packet.

Together, these two techniques greatly reduce excessive preambles, result in the reduction of wasted energy, and allow for lower latency and higher throughput. In addition, both of these techniques are broadly applicable across all forms of digital radios, including packetized and bit stream, because the short time gaps, early acknowledgements, and random backoff can all be implemented in software.

3.4 Packetizing Radios

LPL has a limited ability to support packetizing radios. For example, B-MAC is the default MAC protocol for TinyOS [10] but is incapable of supporting some packet radios such as the Chipcon 2420. B-MAC was originally developed for bit streaming radios like the Chipcon CC1000, which provides low-level access to the individual bits received by the radio. With these radios, B-MAC can generate long preambles. However, the new generation of sensor motes, such as the MicaZ [2], TelosB [4], and iMote [11], make use of the Chipcon CC2420 [1] 802.15.4 radio. Instead of transmitting a raw bit stream, this type of packetizing radio takes as input the payload of the packet, and the radio module inserts its own preamble, header information and CRC. When a packet is received, the radio strips the header, checks the CRC, and if the packet is not corrupted passes the payload of the packet to the microprocessor. While the packet interface is a valuable advance in radio technology, as it reduces the burden on the microprocessor, it limits the ability of the application to precisely control the bits that are sent over the air. Most pertinent, with these radios the application cannot send a preamble of arbitrary length. This precludes the use of LPL protocols that depend on an extended preamble.

For these radios, it is also not possible to mimic an extended preamble by sending a long data packet, which acts as a pseudo-preamble. This is because the receiver will be unable to sample the packet containing the pseudo-preamble, i.e. the packetizing radio will only deliver the packet after it has fully received the entire pseudo-preamble. This defeats the purpose of preamble sampling.

LPL is supported in certain kinds of packetizing radios, such as in the Chipcon 2500 and MaxStream XBee radios, but only because it is implemented directly in the hardware. In this case, long preambles can be specified because the radio supports this configuration option, unlike the Chipcon 2420.

In contrast, X-MAC's short strobed preamble is well-suited to all types of digital radios, as mentioned earlier.

4 Adaptation to Traffic Load

While many sensor network application produce periodic and non-varying traffic, there is also the need to adapt to variable traffic loads. In addition, different nodes in a multi-hop network will experience different average traffic loads. For example, in a network with a tree topology where the nodes sense data periodically and transmit the data over multiple hops to a base station, nodes closer to the base station will receive and transmit more data than those further towards the leaves. As such, nodes in the network must have different sleep schedules to effectively accomodate the different traffic loads. Even for a periodic sensing application, it would be quite difficult for a human operator to hand tune all of the nodes sleep schedules, while for applications with time varying traffic loads hand tuning is simply not possible. Consequently, an effective duty cycling MAC protocol must automatically adapt to varying traffic patterns.

The performance of a duty-cycling MAC is largely determined by the choice of sleep, wake, and radio use periods for both the senders and receivers. One of X-MAC's key contributions, which we describe next, is a lightweight adaptation algorithm which closely approximates the optimal values for these periods.

4.1 Optimality

We consider the following metrics for MAC quality: sender and receiver energy consumption and latency. The achieved throughput equals the offered load at any usage level for which a duty-cycling MAC makes sense. For current devices, the power drawn is almost entirely determined by the node's sleep state and radio operation mode. (See Table 1, appendix A) The expected energy consumption can therefore be modeled in terms of the durations of the sender and receiver sleep, listen, and transmission periods.

We will show that if the probability of receiving a packet in any given interval, $P_d(t)$, is known then sender and receiver tunable parameters can be set to optimal values. Let P_{Tx}, P_{Rx}, and P_s be the power required to transmit, receive, and sleep, respectively. S_p, S_{al}, and S_d denote the duration of the sender's preamble, acknowledgement listen, and data (packet body) transmission periods. R_l and R_s denote the receiver sleep and listen periods.

The expected energy to send a packet is given by:

$$
\begin{aligned}
E_s =\, & (\text{preamble energy + energy per ACK listen}) \\
& * (\text{expected preamble-listen iterations required}) \\
& + (\text{energy to send packet}) \\
=\, & (P_{Tx}S_p + P_{Rx}S_{al})\left(\frac{1}{\left(\frac{R_l - S_p}{R_l + R_s}\right)}\right) + S_d P_{Tx} \quad (1) \\
=\, & \frac{(P_{Tx}S_p + P_{Rx}S_{al})(R_l + R_s)}{R_l - S_p} + S_d P_{Tx}
\end{aligned}
$$

The expected energy to receive a packet is given by:

$$
\begin{aligned}
E_r =\, & (\text{listen cycle energy + sleep cycle energy}) \\
& * (\text{expected iterations for a preamble to arrive}) \\
& + (\text{energy to send an ACK}) \\
& + (\text{energy to receive packet}) \quad (2) \\
=\, & \frac{P_s R_s + P_{Rx}R_l}{1 - (1 - P_d(t))^{(R_l + R_s)}} + P_{Tx}R_a + R_d P_{Rx}
\end{aligned}
$$

5

The expected latency for a single packet is:

$$Lat = (\text{duration of preamble} + \text{ACK listen})$$
$$* (\text{expected number of iterations required})$$
$$+ (\text{duration to send packet})$$

$$= \left(\frac{1}{\left(\frac{R_l - S_p}{R_l + R_s} \right)} \right) * (S_p + S_{al}) + S_d \qquad (3)$$

$$= \frac{(S_p + S_{al})(R_l + R_s)}{R_l - S_p} + S_d$$

These models lead to the following observations, the derivations of which can be found in appendix B.

THEOREM 4.1. *Energy and latency are both minimized when S_p and S_{al} are set to the lowest values which allow for the preamble to be transmitted and ACK received, respectively.*

For any objective function $f(\cdot)$ which is a function of sender energy, receiver energy, and latency:

THEOREM 4.2. *Optimal receiver sleep and listen times for $\min_{R_s, R_l} f(\cdot)$ depend solely on $P_d(t)$ and device constants.*

For any objective function $f(\cdot)$ consisting of a *convexity-preserving* combination of sender energy, receiver energy, and latency:

THEOREM 4.3. *$\min_{R_s, R_l} f(\cdot)$ can be found by standard convex optimization techniques.*

Thus, given an estimate of $P_d(t)$, the protocol parameters can be determined mechanically.

4.2 Approximation

Nonlinear minimization is too demanding a process to be desirable in a sensor networking MAC. That said, the mapping $P_d(t) \to (R_s{}^*, R_l{}^*)$ is smooth enough to admit lightweight approximations. Figure 2 shows this mapping over a range of packet arrival rates.
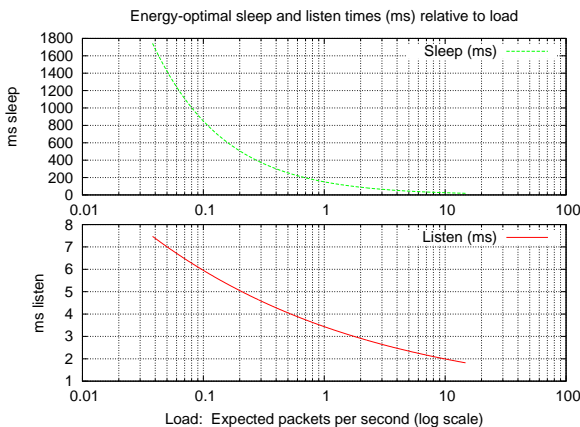

Figure 2. Energy-optimal sleep and listen times

The on-node approximation is based on linear interpolation: We pre-compute a table of $P_d(t)$ values and their associated optimal $R_s{}^*, R_l{}^*$ values[6]. In on-line operation, the sensor node uses its estimate $\widehat{P_d(t)}$ to perform a table lookup and interpolates between the closest pre-computed values.

Numerical simulations suggest that this approximation achieves energy efficiency comparable to direct optimization. We chose an interpolation table of 24 exponentially-spaced entries, ranging between 10^{-4} and 10^3 expected packets per second. The energy-efficiency of the optimal and interpolated values of R_s and R_l were then compared for ten thousand values of P_d. Figures 3 and 4 show the results of this experiment: Figure 3 shows the "raw" difference between optimal and interpolated results, and Figure 4 shows the difference as a fraction of the optimal value. The mean difference is 0.45%, and 95th percentile difference is 1.3%.

Figure 3. Energy waste per packet due to interpolation.
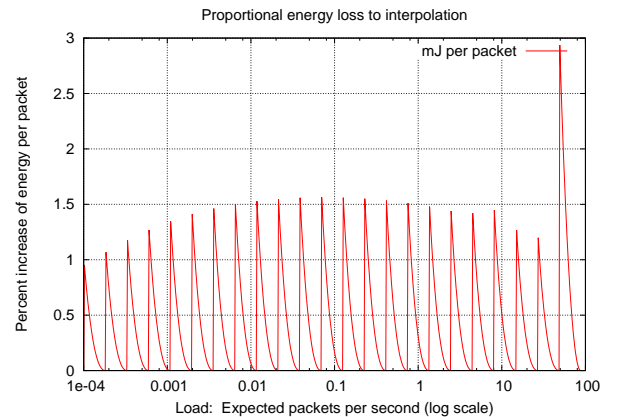
Figure 4. Energy waste as a fraction of optimum.

4.3 Estimating Traffic Load

The preceding produces near-optimal values when the traffic load $P_d(t)$ is known. An estimate of the instantaneous value, $\widehat{P_d(t)}$ can be derived from the observed packet arrival rate: The likelihood of k packets arriving over a period of $n * t$ can be modelled as a Bernoulli process of n trials with probability of success $P_d(t)$. The most likely value of $P_d(t)$ is that which maximizes the probability of the observed outcome. Applying Bayes' rule, the most likely value of $P_d(t)$

is the maximum on the interval $(0, 1)$ of the function:

$$f(P_d | k, n) = \frac{\frac{n!}{(n-k)!k!} P_d{}^k (1 - Pd)^{n-k} f(P_d)}{\int_0^1 \frac{n!}{(n-k)!k!} P_d{}^k (1 - Pd)^{n-k} f(P_d) \, dP_d} \quad (4)$$

THEOREM 4.4. $\widehat{P_d(t)} = \frac{k}{n}$ is an optimal instantaneous estimate of $P_d(t)$.

In the case where there is no prior knowledge of the probability distribution of $P_d(t)$, equation 18 has its maximum at $\frac{k}{n}$. The derivation is given in appendix B.

A moving estimate can be maintained by any of the standard mechanisms; without knowing the dynamics of application load change it is impossible to identify an optimum.

5 Evaluation

In order to evaluate and demonstrate the correctness and benefits of X-MAC we have implemented the protocol on top of the Mantis Operating System (MOS) [5]. MOS is an open source, multi-threaded operating system developed at the University of Colorado at Boulder for use on wireless sensor networking platforms. It is currently ported to the Mica2, MicaZ, and TelosB sensor nodes and is written entirely in C. X-MAC is implemented as a MAC layer protocol and an application initiates the protocol by calling wor_init(), which takes 3 parameters. The parameters are minimum sleep time, maximum sleep time, and default sleep time, which determines the initial sleep period. The minimum sleep time parameter allows the application developer to specify the minimum sleep time that enables the network to achieve some application specific minimum network lifetime. The maximum sleep time parameter allows the developer to bound the one hop latency in order to meet latency requirements.

The wake-on-radio initialization function *wor_init()* spawns a receive thread that queues incoming packets for processing by the application layer, and also wakes and sleeps the radio according to the duty cycle. A second thread, the application thread, is written by the developer and makes use of *wor_send()*, which implements the sending functionality of X-MAC, and the MOS function *com_recv()*, which processes the packets in the receive queue. The application developer is responsible for waking and sleeping the application thread if they wish to completely sleep the sensor node. For these evaluations, the awake time is always 15 ms. This is the amount of time necessary to receive a preamble packet, determine if the node is the intended receiver, and return the acknowledgement packet to the transmitting node.

5.1 Experimental Setup

For our experiments, we used a deployed indoor testbed of TelosB motes. The TelosB platform was developed at the University of California at Berkeley and is marketed and sold by Moteiv and Crossbow. The radio used by the TelosB is the Chipcon CC2420, which is an 802.15.4 compliant device, has a data rate of 250kbps, and operates in the 2.4 GHz ISM band. The mote uses an 8 MHz TI MSP430 and has 1 MB of external flash. The current draw of the device, excluding the radio, is 1.8 mA in active mode and 5.1 μA when in sleep mode. The CC2420 radio consumes 23 mA in receive mode, 17.5 mA when transmitting at 0 dBm, 21μA in idle mode,

and 1 μA in sleep mode. When X-MAC "sleeps" the radio, in fact it puts the radio into idle mode, as sleep mode turns off the oscillator and requires a longer time to transition back to receive mode.

As the energy draw of the radio when receiving is far greater than when in idle mode, it is of the utmost importance to achieve a low duty cycle to extend the life of the network. Because the processor consumes an extremely small amount of energy in comparison with the radio, in our evaluation we allow the application thread to run continuously while the radio is turned on and off according to the duty cycle. If a node wishes to transmit a packet, it turns the radio on and attempts to send the packet. Consequently, the total time that the radio is on is the time spent sampling the channel for preamble packets, the time awake receiving premble packets and data packets, and the time that the radio is awake attempting to send a packet.

As a comparison protocol for X-MAC, we have implemented a simple asynchronous LPL MAC protocol. This protocol is the closest approximation that we could develop using a packetizing radio. When sending, the transmitter sends a stream of preamble packets as rapidly as possible, and after the extended preamble the data packet is sent. There are two differences between X-MAC and the simple LPL protocol; first, the simple protocol does not inspect the preamble packets for the target ID so all receivers will remain awake until they receive the data packet; second, with the simple protocol, transmitters always send the entire extended preamble and receivers do not send an acknowledgement packet to the transmitter. In addition, the adaptation algorithm cannot be applied to the simple protocol. Although the receiver can adjust its sleep period, the transmitter will not be aware of this change so it will not know to adjust the length of its preamble. In X-MAC, the preamble is truncated by the acknowledgement so the transmitter does not need to be explicitly informed of the change in receiver sleep period. When comparing X-MAC to the simple protocol, we attempt to calculate the sleep period that best accomodates the traffic load.

5.2 X-MAC Performance

To show the performance benefits of the X-MAC protocol, we performed a number of simple experiments that evaluate X-MAC without adaptation for simple topologies with no contention and low injection rates. This allows us to see clearly the functioning of X-MAC.

5.2.1 Duty Cycle

To demonstrate the benefits of the overhearing avoidance and the short strobed preamble in X-MAC, we performed an experiment with a varying number of nodes. For this, we set up a star topology where multiple senders are transmitting to a single receiver, and all nodes are within transmission range of each other. Each node sends a packet once every 5 seconds to the receiving node and all nodes have a sleep period, and preamble length, of 500 ms. We proceeded to measure the percentage of time the nodes are in sleep mode. The experiment is repeated with a varying number of senders, from 1 to 5, and the transmissions are timed so as to avoid contention. This means that for the single sender case, there is
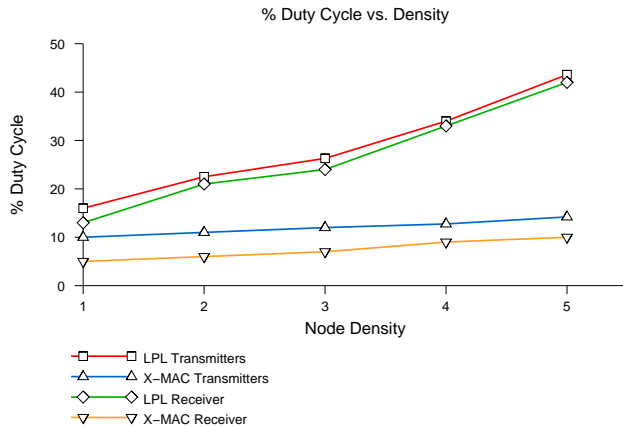
Figure 5. Duty cycles of senders and the receiver as a function of network density.

one packet being sent every 5 seconds while for the 5 sender case there are 5 packets being sent every 5 seconds.

As can be seen in Figure 5, with X-MAC the duty cycle does not vary greatly as network density increases. As there is a small amount of processing time needed to check the ID in the preamble, we see a slight decrease in sleep time at higher densities. In contrast, with the simple LPL asynchronous protocol, nodes wake up when they detect preamble packets transmitted from the other senders, as they do not know if the transmission is intended for them or not. With the simple protocol, as network density increases, the time a node sleeps decreases in proportion to the number of transmitters it can overhear.

For this experiment, we also show the duty cycle of the receiver seperately. This allows us to see the change in receiver duty cycle as the traffic load increases. As can be seen, the receiver is able to sleep more when using X-MAC because it sends the early acknowledgement and returns to sleep. The ability of X-MAC to substantially prolong energy lifetime is highlighted by looking at the single sender case. In this case, the receiver's duty cycle is 5% under X-MAC and 13% under LPL, resulting in a doubling to tripling of the energy lifetime of the receiver for X-MAC.

Additionally, X-MAC further reduces energy consumption at the transmitter by having the receiver send an acknowledgement packet when it awakens, effectively truncating the preamble. In X-MAC, when a preamble is detected the receiver sends the early acknowledgement packet and the transmitter immediately sends the data packet and both nodes return to sleep; thus saving energy at both the sender and the receiver. In contrast, with the simple protocol the receiver must remain awake for the remainder of the preamble and the transmitter sends the entire extended preamble before every transmission, thus wasting energy. The benefit of the shortened preamble is highlighted by looking at the duty cycle of the single sender case for X-MAC and the simple protocol. Here, it can be seen that, while both protocols send one packet with one preamble, X-MAC uses substantially less energy because its preamble is shortened by the
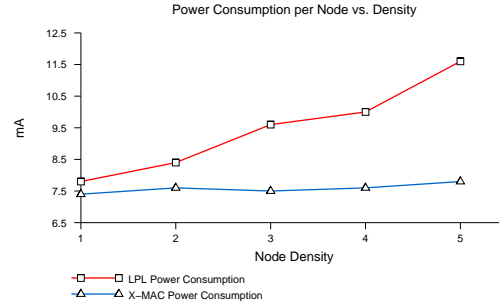
Figure 6. Power consumption per node versus density.

Figure 7. Latency

receiver, i.e. the duty cycle is 10% for X-MAC versus 17% for LPL.

5.2.2 Energy Usage

In order to show the energy savings in terms of actual power consumed, we attached an oscilloscope to one of the transmitting nodes, and repeated the above experiment. We measured the mean current draw in mA. As can be seen in Figure 6, the energy consumed by the simple protocol increases as network density increases. For X-MAC, energy consumption remains constant as network density increases. These results confirm our previous measurements of the duty cycles.

5.2.3 Latency

In X-MAC, the extended preamble is truncated when the receiver wakes up. While this provides significant energy savings, it also reduces per hop latency because the receiving node can immediately begin forwarding the packet after it receives it. To show the reduction in latency when using X-MAC, we use a chain topology of 5 nodes. The sleep period of all nodes is 500 ms, as is the preamble length. We generate one packet every 5 seconds at one end of the chain, transmit the packet over 4 hops to the far end of the chain, and then forward the packet in the reverse direction back to the originating node. We then measure the round trip time of the packet. This reverse path is necessary to be sure that the same clock that is used to time stamp the packet initially is again used to measure the round trip time.

In Figure 7, we show the results of our 8 hop latency test. For the simple protocol, the round trip time is nearly 4 seconds, as would be expected because the packet incurs 500 ms of delay at each of 8 hops. In contrast, X-MAC experiences a round trip time closer to 2.5 seconds. Analytically, the round trip time should be about 2 seconds, as the receiver

will wake up, on average, half way through the preamble. Of course, the receiver does not always wake up exactly half way through the preamble. Our experiments show that X-MAC reduces latency by approximately 50%.

6 Discussion and Future Work

While our experiments have highlighted the substantial first-order performance gains of X-MAC, more detailed experiments are needed. Our duty cycle measurements are in need of error bars. Our latency measurement needs to be confirmed with more data points. We also would like to vary other parameters, such as packet length, data rate, etc. We intend to evaluate X-MAC on a tree-structured routing topology. We also plan to evaluate X-MAC on a larger 50-node testbed of motes. One of the key properties of MAC protocol design is fairness, and we hope to evaluate the fairness of X-MAC in more detail. We were unable to evelute the adaptation algorithm and intend to compare its performance and stability properties against basic B-MAC.

We estimate the memory consumption of the full X-MAC implementation to consume 6-7 KB of flash. We did not attempt to optimize this implementation at all, and instead sought ease of implementation. We intend to explore a more compact code implementation in the future. We also intend to provide RAM consumption numbers in the next iteration. For example, the table used for approximating the optimal adaptation algorithm consists of only 20 integers stored in RAM, but the other RAM consumption numbers need to be determined.

X-MAC consumes a minimum of 15 ms listen time. We need to understand in more detail the various contributions to this latency, so that we can reduce this minimum listen time further.

We hope to have an implementation of X-MAC for TinyOS in the near future. Since X-MAC does not depend upon the radio being either bit-oriented or packet-oriented, then X-MAC should enable TinyOS to operate our new approach to LPL universally across MICA2, MICAz, and TelosB motes.

7 Conclusions

This paper describes X-MAC, a new approach to low power communication in WSNs. X-MAC employs a strobed preamble approach by transmitting a series of short preamble packets, each containing the address information of the target receiver. The series of short preamble packets approximates a continuous preamble. Small pauses between preamble packets permit the target receiver to send an acknowledgment that stops the stream of preamble packets, thereby truncating the extended preamble. The advantages of this approach are multi-fold: the transmitter does not need to send the full extended preamble, thus saving energy at both the transmitter and receiver and allowing for lower per hop latency and higher throughput; non-target receivers who overhear the strobed preamble can go back to sleep immediately, as compared to remaining awake for the full preamble as in conventional LPL; and, this strobed preamble approach can be readily adapted to the packetized radios that are emerging as the standard in today's sensor motes. Another key feature of X-MAC is its algorithm for adapting the duty cycle of the receiver automatically to adapt to varying traffic loads. We verified that X-MAC's strobed preamble approach outperforms traditional LPL by implementing the protocol and performing an array of experiments using a testbed of wireless sensor nodes. X-MAC is seen to lower per-hop latency, reduce energy consumption, and allow for increased throughput.

8 References

[1] Chipcon cc 2500 radios. http://www.chipcon.com.

[2] Crossbow micaz motes. http://www.xbow.com.

[3] Maxstream xbee radios. http://www.maxstream.com.

[4] Moteiv telosb motes. http://www.moteiv.com.

[5] S. Bhatti, J. Carlson, H. Dai, J. Deng, J. Rose, A. Sheth, B. Shucker, C. Gruenwald, A. Torgerson, and R. Han. Mantis os: An embedded multithreaded operating system for wireless micro sensor platforms. *ACM/Kluwer Mobile Networks & Applications (MONET), Special Issue on Wireless Sensor Networks*, 10(4):563–579, August 2005.

[6] J. Dennis and R. Schnabel. Numerical methods for unconstrained optimization and nonlinear equations. Prentice-Hall Series in Computational Mathematics. Prentice-Hall, 1983.

[7] A. El-Hoiydi and J. Decotignie. Low power downlink mac protocols for infrastructure wireless sensor networks. *ACM Mobile Networks and Applications*, 10(5):675–690, 2005.

[8] L. Feeney and M. Nilsson. Investigating the energy consumption of a wireless network interface in an ad hoc networking environment. In *IEEE INFOCOM*, volume 3, pages 1548–1557, 2001.

[9] G. Halkes, T. V. Dam, and K. Langendoen. Comparing energy-saving mac protocols for wireless sensor networks. *ACM Mobile Networks and Applications*, 10(5):783–791, 2005.

[10] J. Hill, R. Szewczyk, A. Woo, S. Hollar, D. Culler, and K. Pister. System architecture directions for network sensors. In *Nineth International Conference on Architectural Support for Programming Languages and Operating Systems (ASPLOS'00)*, pages 93–104, Cambridge, MA, USA, November 2000.

[11] L. Krishnamurthy, R. Adler, P. Buonadonna, J. Chhabra, M. Flanigan, N. Kushalnagar, L. Nachman, and M. Yarvis. Design and deployment of industrial sensor networks: experiences from a semiconductor plant and the north sea. In *SenSys '05: Proceedings of the 3rd international conference on Embedded networked sensor systems*, pages 64–75, 2005.

[12] N. Mishra, K. Chebrolu, B. Raman, and A. Pathak. Wake-on-wlan. In *The 15th Annual Interntional World Wide Web Conference (WWW), to appear*, 2006.

[13] B. O'Hara and A. Petrick. *The IEEE 802.11 Handbook: A Designer's Companion*. Institute of Electrical

& Electronics Engineers, 2nd edition, 2005.

[14] J. Polastre, J. Hill, and D. Culler. Versatile low power media access for wireless sensor networks. In *The Second ACM Conference on Embedded Networked Sensor Systems (SenSys)*, pages 95–107, November 2004.

[15] E. Shih, P. Bahl, and M. Sinclair. Wake on wireless: An event driven energy saving strategy for battery operated devices. In *MobiCom*, pages 53–64, 2002.

[16] T. van Dam and K. Langendoen. An adaptive energy-efficient mac protocol for wireless sensor networks. In *1st ACM Conference on Embedded Networked Sensor Systems (SenSys)*, pages 171–180, 2003.

[17] W. Ye, J. Heidemann, and D. Estrin. An energy-efficient mac protocol for wireless sensor networks. In *21st International Annual Joint Conference of the IEEE Computer and Communications Societies (IN-FOCOM'02)*, New York, NY, USA 2002.

[18] W. Ye, J. Heidemann, and D. Estrin. Medium access control with coordinated, adaptive sleeping for wireless sensor networks. *ACM Transactions on Networking*, 12(3):493–506, June 2004.

A Definitions
A.1 Variables

$$P_{Tx} \triangleq \text{Power to Tx}$$

$$P_s \triangleq \text{Power to sleep}$$

$$P_{Rx} \triangleq \text{Power to Rx}$$

$$S_p \triangleq \text{Duration of sender preamble}$$

$$S_{al} \triangleq \text{Duration of sender ACK listen}$$

$$S_d \triangleq \text{Duration of sender data Tx}$$

$$R_s \triangleq \text{Duration of receiver sleep}$$

$$R_a \triangleq \text{Duration of ACK send}$$

$$R_l \triangleq \text{Duration of receiver listen}$$

$$R_d \triangleq \text{Duration of receiver data Rx}(= S_d)$$

$$P_d(t) \triangleq \text{Packet probability per time } t$$

A.2 Constraints

The following constraints describe the range of variable values for which our model is reasonable:

$$\{P_{Tx}, P_s, P_{Rx}, S_p, S_{al}, S_d, R_s, R_a, R_l, R_d\} > 0 \tag{5}$$

$$0 \le P_d(t) \le 1 \qquad \text{probability range} \tag{6}$$

$$R_l > S_p \qquad \text{preamble reception possible} \tag{7}$$

$$S_{al} > R_a \qquad \text{ACK reception possible} \tag{8}$$

$$R_d = S_d \qquad \text{No Doppler effect} \tag{9}$$

A.3 Concrete Equations

Table 1 gives device-specific constants for the Telos IV mote.

Variable	value	source
P_{Tx}	57.6 mW	device-specific
P_s	0.0183 mW	device-specific
P_{Rx}	74.4 mW	device-specific
S_p	0.26 ms	(msg. size)/bw + overhead
R_a	0.26 ms	(msg. size)/bw + overhead
S_d, R_d		application-specific
$P_d \triangleq P_d(1ms)$		application-specific
S_{al}	0.26 ms	free, but $S_{al} = R_a$ is optimal.
R_s	$0 \le R_s$	
R_l	$.26 \text{ ms} \le R_l$	

Table 1. Variable values

Substituting these values into equations 1 - 3 gives the following, where time, power, and energy are measured in ms, mW and μJ respectively:

$$E_s = 57.6\,S_d + \frac{34.32\,(R_s + R_l)}{R_l - 0.26} \tag{10}$$

$$E_e = \frac{0.0183\,R_s + 74.4\,R_l}{1 - (1 - P_d)^{R_s + R_l}} + 74.4\,S_d + 14.976 \tag{11}$$

$$Lat = S_d + \frac{0.52\,(R_s + R_l)}{R_l - 0.26} \tag{12}$$

B Proofs
B.1 Theorem 4.1
B.1.1 Sender Preamble Time

The sender preamble duration S_p affects the expected energy to send and the expected latency (equations 1 and 3.) Both take their optimal (minimal) values when S_p is minimized, as will be shown shortly. S_p is bounded from below by the message size / the available bandwidth + processing overhead. The partial derivative of equations 1 and 3 with respect to S_p are given below.

$$\frac{\partial E_s}{\partial S_p} =$$

$$\frac{(R_s + R_l)\,(P_{Tx}S_p + P_{Rx}S_{al})}{(R_l - S_p)^2} + \frac{P_{Tx}\,(R_s + R_l)}{R_l - S_p} \tag{13}$$

$$\frac{\partial Lat}{\partial S_p} = \frac{(R_s + R_l)\,(S_p + S_{al})}{(R_l - S_p)^2} + \frac{R_s + R_l}{R_l - S_p} \tag{14}$$

It follows from constraints 5 and 7 that $\frac{\partial E_s}{\partial S_p} > 0$ and $\frac{\partial Lat}{\partial S_p} > 0$ for all permissible values. Thus, within the defined bounds, the expected latency and expected sender energy consumption are always reduced by lowering S_p, and expected receiver energy consumption is unaffected.

B.1.2 Sender ACK Listen Time

The sender acknowledge listen time, S_{al} also affects equations 1 and 3.

The partial derivative of equation 1 with respect to S_{al} is given below.

$$\frac{\partial E_s}{\partial S_{al}} = \frac{P_{Rx}(R_s + R_l)}{R_l - S_p} \qquad (15)$$

It follows from constraints 5 and 7 that $\frac{\partial E_s}{\partial S_{al}} \geq 0$ for all feasible values of all variables. Consequently, E_s is always minimized when the lowest permissible value is chosen for S_{al}.

Similarly for the expected latency,

$$\frac{\partial Lat}{\partial S_{al}} = \frac{R_s + R_l}{R_l - S_p} \qquad (16)$$

it is always the case that $\frac{\partial Lat}{\partial S_{al}} \geq 0$.

From the preceding paragraphs, it follows that latency, receiver energy consumption and sender energy consumption all take minimal values when S_{al} is minimized.

B.2 Theorem 4.2

Once the device attributes and the parameters with invariant optimal values are fixed, all three objective functions are given by equations 10 - 12. The values of the objectives depend on, at most, $\{P_d, S_d, R_l, R_s\}$.

P_d and S_d can be regarded as unknown constants[1]. Thus, given any particular P_d and S_d and a constrained range values for R_l and R_s, there exist minimal values E_s^*, E_r^*, and Lat^*.

For any given (P_d, S_d), for each objective or for any combination thereof, there exists a non-empty set of (R_l^*, R_s^*) pairs producing the minimal value of the objective. Additionally, S_d does not appear in the partial derivative of E_s, E_r, or Lat with respect to R_l or R_s. Consequently, the sets of values which minimize those objectives do not depend on the value of S_d. Thus, any objective based on some combination of E_s, E_r, and Lat, can be be written as some $f(P_d, R_l, R_s) : \mathbb{R}^3 \to \mathbb{R}$.

B.3 Theorem 4.3

Consider $f(x) \in \{E_s, E_r, Lat\}$. Each function is convex over the domain x consistent with the constraints given in appendix A.2. All three functions are twice differentiable over this domain, and thus have a well-defined Hessian matrix. For each f, $\forall x \in \mathbf{dom} f$:

$$\triangledown^2 f(x) \succeq 0$$

Therefore, each f is convex over the appropriate range of R_s and R_l. By extention, any $g(x)$ which consists of convexity-preserving combinations of these fs is also convex. A notable group of these is the set of nonnegative linear combinations of fs.

Any local minimum of $g(x)$ within the appropriate range will therefore also be a global minimum, which means that many non-linear programming techniques can be applied.

---

[1]They can, of course, change value over time.

B.4 Theorem 4.4

For any given $P_d(t)$, the probability that k packets will arrive over duration nt can be modelled as a Bernoulli process of n trials with probability of success Pd. The frequency mass function for the probability of k "hits" in n trials is given by

$$P_n(k) = f(k, n | P_d) = \binom{n}{k} P_d{}^k (1 - P_d)^{n-k} \qquad (17)$$

Applying Bayes' rule, we get the following frequency density function:

$$\begin{aligned} f(P_d | k, n) &= \frac{f(k, n | P_d) f(P_d)}{\int_0^1 f(k, t | P_d) f(P_d) \, dP_d} \\ &= \frac{\frac{n!}{(n-k)!k!} P_d{}^k (1 - Pd)^{n-k} f(P_d)}{\int_0^1 \frac{n!}{(n-k)!k!} P_d{}^k (1 - Pd)^{n-k} f(P_d) \, dP_d} \end{aligned} \qquad (18)$$

In the case where there is no prior information about the distribution, that is where $f(P_d)$ is uniform, and $k \geq 0$, equation 18 reduces to:

$$f(P_d | k, n) = \frac{(1 - P_d)^{n-k} P_d{}^k}{\int_0^1 (1 - P_d)^{n-k} P_d{}^k \, dP_d} \qquad (19)$$

The best estimate $\widehat{P_d(t)}$ is the value of P_d which maximizes $f(P_d | k, n)$. Note that equations 18 and 19 are undefined where P_d is 0 or 1. For the special cases $k = 0$ and $k = n$, $f(P_d | k, n)$ has no extrema in the interval [0,1]. Where $k = 0$, the maximum is found where $P_d = 0$ and similarly where $k = n$, the maximum occurs where $P_d = 1$. For $0 < k < n$, the following analysis holds:

$$\frac{df}{dP_d} = \frac{k (1 - P_d)^{n-k} P_d{}^{k-1}}{\int_0^1 (1 - P_d)^{n-k} P_d{}^k \, dP_d} - \frac{(n-k) (1 - P_d)^{n-k-1} P_d{}^k}{\int_0^1 (1 - P_d)^{n-k} P_d{}^k \, dP_d} \qquad (20)$$

$\frac{df}{dP_d}$ is zero where:

$$P_d{}^k = -\frac{k (1 - P_d) P_d{}^{k-1}}{k - n} \qquad (21)$$

Equation 21 has two solutions: $P_d = 0$ and $P_d = \frac{k}{n}$. As mentioned above, equations 18 and 19 are undefined where P_d is 0. Thus:

$$\widehat{P_d(t)} = \begin{cases} 0 & \text{if } k = 0 \\ \frac{k}{n} & \text{if } 0 < k < 1 \\ 1 & \text{if } k = n \end{cases} \qquad (22)$$

This is of course just $\widehat{P_d(t)} = \frac{k}{n}$ for $0 \leq k \leq n$.