

An Adaptive Link Layer for Range Diversity in Multi-radio Mobile Sensor Networks

Jeremy Gummesson, Deepak Ganesan, Mark D. Corner and Prashant Shenoy

Department of Computer Science
University of Massachusetts, Amherst
Amherst, MA 01003

Email: {gummesson, dganesan, mcorner, shenoy}@cs.umass.edu

Abstract—An important challenge in mobile sensor networks is to enable energy-efficient communication over a diversity of distances while being robust to wireless effects caused by node mobility. In this paper, we argue that the pairing of two complementary radios with heterogeneous range characteristics enables greater range diversity at lower energy cost than a single radio. We make three contributions towards the design of such multi-radio mobile sensor systems. First, we present the design of a novel reinforcement learning-based link layer algorithm that continually learns channel characteristics and dynamically decides when to switch between radios. Second, we describe a simple protocol that translates the benefits of the adaptive link layer into practice in an energy-efficient manner. Third, we present the design of *Arthropod*, a mote-class sensor platform that combines two such heterogeneous radios (XE1205 and CC2420) and our implementation of the Q-learning based switching protocol in TinyOS 2.0. Using experiments conducted in a variety of urban and forested environments, we show that our system achieves up to 52% energy gains over a single radio system.

I. INTRODUCTION

The choice of the wireless radio is perhaps the single most crucial design parameter for designing mobile sensor networks such as DieselNet [2] or ZebraNet [8]¹. The wireless radio must enable node-to-node and node-to-basestation communication over distances dictated by application needs, while being energy-efficient and robust to wireless effects introduced by mobility patterns. With advances in communication technologies, a spectrum of wireless radios are available to meet the needs of a sensor network. Table I depicts four common wireless radios used by today’s sensor network platforms. As shown in the table, wireless radios are generally designed with a communication range in mind. For example, the Xtend and the XE1205 radios are designed for low-bitrate long-range communication over distances of a mile or more. In contrast, 802.11 and CC2420 radios enable high and low bandwidth communication, respectively, over short ranges of hundreds of feet or less. Thus, the sensor network designer must make a critical design choice. She can either choose a long-range radio enabling nodes to communicate over long distances but at the expense of expending more power. Or she can choose a shorter range radio that is more power-efficient but forego communication over longer distances.

Note that traditional techniques for range adaptation via power control or range elongation via the use of directional antennas do not address this tradeoff for mobile sensor networks. As shown in Table I, modern radios support range adaptation using power control — a higher power setting can be used to increase the communication range of the radio. While it is possible to choose a long range radio and use lower power settings for short range communication, doing so is far less efficient than using a short range radio for communicating over shorter distances. As shown in Table I, the lowest power setting on the XTend radio is still 561x more expensive than using the CC2420 radio. Similarly, it is not feasible to use a radio designed for short range communication and to “increase” its range by using directional antennas. Directional antennas have been used successfully to increase the communication range of such radios – for example, the Mobisteer project [11]. However, since directional antennas are bulky, it is not feasible to deploy them in many mobile sensor network settings; for instance, animal tracking deployments require compact packaging of the mobile sensors.

In this paper, we argue that pairing two complementary radios with heterogeneous range characteristics enables mobile sensor nodes to achieve a significantly greater range diversity at a lower total energy cost when compared to a single radio. The key idea is to operate each radio over a range where it is more efficient and to switch to the other radio whenever a mobile node moves from one radio’s effective range to another. In this manner, we achieve the best-of-both-worlds and eliminate the drawbacks of a single radio platform.

We present the design of a heterogeneous multi-radio platform and system for handling range dynamics, where the choice of which radio to use for communication is made dynamically based on current channel characteristics, specifically wireless channel variations caused by device mobility and range effects. To shield applications from the increased complexity of choosing between radios, we present the design of a unified link layer that *transparently* chooses which radio to employ for communication between a pair of nodes. At the core of such a link layer is an adaptive algorithm that can dynamically decide *when* to use each radio for a wide range of mobility patterns. Such an algorithm is non-trivial since it needs to continually monitor and “learn” channel characteristics for the two radios and determine which one provides

¹This research was supported by NSF grants CNS-0626873, CNS-0615075, CNS-0520729, CNS-0546177, and EEC-0313747.

Radio	Bandwidth	transmit power levels (min, max), steps	transmit energy/bit (min,max)	max outdoor range
CC2420	250 Kbps	(-25,0dBm),31	102,208nJ/bit	80m
802.11b	11 Mbps	(0,15dBm),4	-,120nJ/bit	100m
XE1205	38.1 Kbps	(0,15dBm),4	1803,5276nJ/bit	80m - 800m
XTend	9.6kbps	(0,30dBm),4	57.3,380.2uJ/bit	2-3km

TABLE I
A SPECTRUM OF RADIO HARDWARE

the lowest energy communication channel. Additionally, the practical implementation of such an adaptive link layer on sensor platforms presents a significant challenge since the energy and resource overhead for monitoring, learning, and switching between radios needs to be kept as low as possible.

Contributions

In this paper, we propose a novel multi-radio hardware and link layer solution for range-adaptive mobile wireless sensor networks. Our work has three major contributions:

Q-Learning based Unified Link Layer: Our first contribution is a reinforcement-learning based algorithm that enables adaptation across radios with different power/range tradeoffs. This algorithm learns the characteristics of radio channels through exploration and continually adapts to use the more efficient one.

Multi-radio Switching Protocol: Our second contribution is a energy-efficient switching protocol that translates the benefits of the Q-learning based adaptation algorithm into practice. The protocol transparently switches between radios, thereby providing the abstraction of a unified link layer to applications executing on multi-radio platforms.

Heterogeneous Multi-Radio Sensor Platform: Our third contribution is the design of a new mote-class sensor platform, the *Arthropod*, that pairs two radios with complementary characteristics: the CC2420 and XE1205. These radios have very different maximum ranges (80 meters vs 800 meters), and also significantly differ in their maximum power output (0 dBm vs 15 dBm) Thus, the *Arthropod* offers good potential for range adaptation to handle mobility effects.

We conduct mobility experiments using our hardware and software prototype in a variety of settings—urban/indoor, urban/outdoor, foliage—and for a range of mobility patterns—continuous and nomadic—that are typical in mobile sensor network deployments. Our experiments show that we obtain up to 52% improvements in energy efficiency over using only one of the two radios on the platform, while achieving a loss rate only marginally higher than using just the high-power radio.

II. RELATED WORK

Since radio diversity presents clear benefits along a number of dimensions: energy, robustness to interference, increased bandwidth and ease of deployment, a number of multi-radio systems have been designed in recent years. This has primarily involved a separation of control tasks such as neighbor discovery or neighbor wakeup from data transmission. Such

a separation has been achieved by pairing 802.11 with the CC2420 [10] or the CC1000 [9], [13], [18], 802.11 with a custom radio for Wake-On-Wireless [17], and 802.11 with an XTend [21] radio [2] for the UMassDieselNet DTN [4]. While such static allocation of roles to radios offers useful benefits, it does not fully utilize the potential of multi-radio systems. In our system, either radio can be used for control or data communication and the choice of which radio to use for communication is made dynamically based on current channel characteristics.

Multiple radio interfaces have also been exploited for increasing bandwidth and tolerating disconnection on mobile wireless devices. The Mobile Access Router [16] exploits multiple types of radio interfaces (eg. 802.11, GPRS, etc), or interfaces tied to different service providers to aggregate bandwidth and avoid stalled transfers. A related technique is PTCP that uses link-layer striping [7] to achieve a similar goal. All these mechanisms are aggressive in using multiple interfaces and do not take energy into account when choosing an interface. An updated Wake-On-Wireless system [1] and Context-for-Wireless [15] use 802.11 with cellular radios for data transmissions, with a static preference given to 802.11 when available.

One dynamic, energy-aware system is Coolspots, which combines 802.11 with Bluetooth [12]. Coolspots chooses Bluetooth transmission when available, and 802.11 when Bluetooth is insufficient to meet the bandwidth requirements. However, the choice of when to use a radio is made using coarse-grained feedback from the network layer, and neglects the benefits of a fine-grained, link-layer approach; this type of approach is useful because it allows a system to react quickly to short term dynamics. Other systems, such as Triage extend this paradigm from multiple radios to multiple platforms [3]; however in this work a single platform is sufficient to process data transmission from both radios.

Recent work on wireless mesh networks has explored designs with multiple radios per node. For instance, carefully planned mesh networks can exploit multiple radios to make channel assignment more effective [5]. However, these approaches have not addressed the problem of algorithms to dynamically react to changing channel characteristics, and do not consider energy efficiency.

In sum, it is our view that ours is the first system to use multiple low-power radios for link-layer, energy-aware, transmission of data, and constitutes an important step towards a truly fine-grained, energy-adaptive, multi-radio link layer system.

Algorithm 1 Q-Learning

- 1: Initialize $Q(s, a)$ arbitrarily
- 2: Repeat(for each step of episode):
- 3: Choose a from s using policy derived from Q (ϵ -greedy)
- 4: Take action a , observe r, s'
- 5: $Q(s, a) \leftarrow Q(s, a) + \alpha[r + \gamma \max_{a'} Q(s', a') - Q(s, a)]$
- 6: $s \leftarrow s'$;
- 7: until s is terminal

III. Q-LEARNING BASED MULTI-RADIO LINK LAYER

Even more so than their fixed counterparts, radios in mobile environments incur unpredictable and widely varying conditions due to channel effects such as shadowing, fading, and multi-path effects, as well as varying interference. While adapting using multiple radios, the channel has hidden state: conditions on the radio not being used. In order to avoid local minimum, the system must periodically attempt to explore other operating states.

Such a system can be compactly implemented using learning algorithms. In particular, we have chosen to use a reinforcement technique called Q-Learning that provides exactly the properties required: a simple reward for making correct decisions and an ability to explore other operating points periodically [19]. In this section, we introduce some concepts from reinforcement learning and outline the design of the adaptation algorithm that is at the core of our unified multi-radio link layer.

A. Introduction to Q-Learning

Q-Learning is a reinforcement-learning technique to enable decision-making for agents in an unknown environment [19]. An agent continually takes an action from a set of possible actions and observes some reward associated with the outcome of their decision. In Q-Learning, there is a "Q-Matrix" that updates according to the reward received, and the state transitions. This Q-Matrix is used to determine which action is optimum while an agent is in a given state. A Q-Matrix has a corresponding Reward-Matrix that contains the reward to be received by the agent for arriving at a particular state. In Q-Learning, the agent will occasionally take a random action to explore states that have not been visited for some time.

Algorithm 1 shows the procedure for Q-Learning [19]. The learning algorithm uses three parameters — the learning rate α , discount factor γ , and the ϵ -greedy parameter. The learning rate places a limit on how quickly learning occurs. If this parameter is set too low, it will take a long time for the system to learn, while if set too high, will cause the Q-Values to never converge to optimal values. The discount factor is used to determine how much emphasis is placed on future rewards. Setting this parameter low will optimize for immediate rewards, while setting this parameter high will place more importance on future rewards. Parameter ϵ determines with what probability a random action is explored, rather than choosing an action with the highest Q-Value.

B. Designing a Unified Link Layer using Q-Learning

We now describe how Q-Learning can be used to adapt between different radios in the case of a dual-radio sensor node. Our discussion assumes the availability of a multi-radio platform with complementary radios (e.g. a short-range low-power and long-range high-power radio). In Section V, we describe one such platform that we have designed that combines a CC2420 and XE1205 radios. We also assume a typical DTN traffic model where mobile nodes buffer data in bundles during periods of disconnection, and transfer data between each other when they come in contact. For example, in DieselNet [2], when a bus node comes in contact with a stationary roadside AP, the bus node first transfers bundles that it has to the AP, followed by the AP transferring buffered bundles to the bus.

We first consider the case where each radio is set to a single power level. In this case, Q-Learning uses a two state model (one for each radio) where the action that is taken by the agent is either to stay with the same radio or switch to the alternate radio. The agent will switch radios if the conditions deteriorate on the current radio (or if it gets disconnected), or if conditions improve on the alternate radio. These dynamics are captured in the Q-Matrix at a rate governed by α and γ through feedback from packets transmitted over the current radio, as well as from exploration packets transmitted over the other radio interface (with frequency determined by ϵ). If the agent finds that the alternate radio has a higher Q-Value (*i.e.* lower energy consumption), the agent will choose to use this new radio interface.

This two state model may be expanded to an n -state model, where each state represents a radio at a particular transmit power level, each representing a particular range/power trade-off. For example, four states would be required for two radios, each with two transmit power level options. However, increasing the number of states comes at the cost of either increased exploration overhead or decreased exploration frequency since exploration requires time and energy. We reduce this overhead in the n -state case by considering only three states at a time — the current state and two adjacent states, a lower-range/lower-power state, and a higher-power/higher-range state. Both these adjacent states could be on the same radio or a different radio. Thus, exploration is limited to only two states at any time.

Reward matrix: A key aspect of Q-Learning adaptation is defining the reward matrix R for each state. The unified link layer receives information about the number of retransmission attempts and number of congestion backoffs for each packet that it transmits through either radio; these metrics are used to determine the reward for the current choice of radio / power level. We model the reward as an estimate of the amount of energy associated with the channel metrics collected for a given packet. The amount of energy to transmit a given packet is a function of packet size, static radio parameters such as receive/transmit power and channel sense time, number of retransmission attempts, and the number of congestion backoffs. Energy is a cost, rather than a reward, so its value

is negative. The following equation shows how rewards are calculated where i is the number of retransmissions:

$$r[i] = -(i \cdot (PacketSize \cdot ByteTime \cdot TxPower + AckTimeout \cdot RxPower) + RxPower \cdot (AckRTT) + PacketSize \cdot ByteTime \cdot TxPower)$$

While the above equation determines the reward when a packet is *successfully* transmitted, we also need to consider the case when a packet is unsuccessful after a pre-defined maximum number of retries. In this case, we want the Q-Learning algorithm to progressively try higher power states until it reaches the highest power state. To obtain this behavior, when a packet transmission is unsuccessful on a low-power state, we assign a large negative reward to encourage the algorithm to switch to a higher power state sooner, thereby limiting the number of lost packets. Once the highest power state is reached, if packet transmission is still unsuccessful, a zero reward is assigned since there is no point in switching back to other lower power states until the connection is re-established at the high-power state.

IV. MULTI-RADIO SWITCHING PROTOCOL

Translating the Q-Learning based switching algorithm to a working protocol presents a non-trivial challenge. When a sender decides to switch to or explore another radio, it needs to notify the receiver of such an action. However, such an explicit handoff may not always work, for example, the receiver may be unreachable by the currently used radio due to mobility. A trivial solution would be for the receiver to keep both radios active at all times, obviating the need for handoff. However, this option is clearly inefficient since it requires both radios to be in receive mode, consuming significant energy. Thus, a key challenge that we address is: *how can we design a practical protocol for switching between radios that is energy-efficient and reliable?*

In the rest of this section, we describe the sender and receiver side design for our adaptive multi-radio block transfer protocol. For simplicity, we consider a dual-radio system with a high-power radio (HIGH) and low-power radio (LOW) with only one power level per radio.

A. Sender State Machine

The state machine at the sender is shown in Figure 1. We first describe the normal operation of the state machine before discussing how we handle exceptional cases that arise due to losses and disconnections. When data transfer starts, the sender first needs to “wakeup” the receiver from its IDLE state. There are many approaches to duty-cycling and wakeup (e.g. SMAC [22], BMAC [14]), and we assume that one of these approaches are available for the radio. Once the wakeup command is successful, the sender transitions from IDLE to HIGH-ON state.

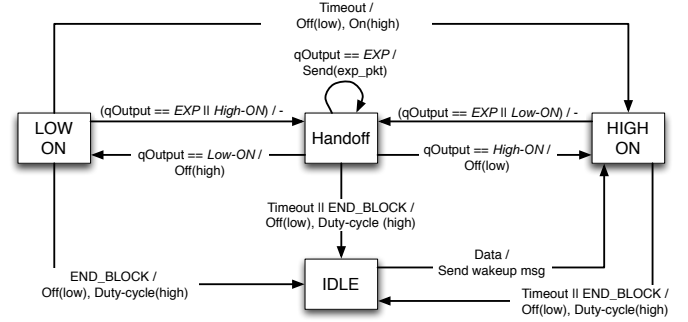


Fig. 1. Sender state machine. qOutput denotes the output of the Q-Learning algorithm, which can be either explore, turn on low-power radio (low), or turn on the high-power radio (high). Transitioning from the IDLE state requires a wakeup message.

Switching and exploration between the radios requires a handshake between the sender and receiver; first, the sender sends a packet indicating that a switch needs to be done, and if the packet is transmitted successfully, the sender and receiver can synchronously switch states to the second radio or explore on it. To perform such a handshake, the sender state machine includes a handoff state in which both radios are turned on. To illustrate, consider a switch from the HIGH-ON to LOW-ON state triggered by the Q-Learning algorithm. The state machine first sends a data packet while remaining in the current state with the *handoff flag* set. If the packet is successfully transmitted, the state machine transitions to the HANDOFF state. (Note that the receiver is in the BOTH-ON state at this point and can receive on both radios). From this state, the sender can send a packet on the LOW radio to transition to LOW-ON state. A similar process is done during exploration. The sender and receiver transition synchronously to the HANDOFF and BOTH-ON states respectively, and stay in this state until exploration is complete, after which they switch back to whatever state they were in earlier.

Finally, we also deal with various cases where the state machines at the sender and receiver may become out-of-sync due to lost packets/acks, or complete loss of connectivity on one or both radios. If the LOW radio is currently in use and becomes disconnected, the sender times out, transitions to the HIGH-ON state and attempts to transmit using the long range radio. (Note that the receiver switches to the BOTH-ON state after a similar timeout, and is ready to receive on the HIGH radio.). If this fails as well, then after another timeout, the sender switches to IDLE mode since it means that the sender and receiver are out-of-range of both radios.

B. Receiver State Machine

The state machine at the receiver is shown in Figure 2. When data transfer starts, the receiver is in the IDLE state, where it operates with the HIGH radio in duty-cycled mode, and the LOW radio in off mode. This enables wakeup by the long-range radio to maximize contact time between the sender and receiver. The receiver is woken up out of this state by a




Fig. 2. Receiver state machine

long preamble on the HIGH radio, and switches to the *HIGH-ON* state. Switching between the two radios occurs through a handoff state where both radios are switched on and ready to receive. When the receiver gets a packet with the *handoff* flag set, it transitions to the BOTH-ON state. It stays in this state until the sender informs the receiver to switch to either the LOW-ON or the HIGH-ON state. The receiver transitions back to the IDLE mode when the END_BLOCK flag is set in a packet indicating that the sender has completed the current transfer of a block.


The receiver state machine also handles a number of exceptional cases that may arise. When the receiver is in the LOW-ON or HIGH-ON state and does not receive a packet for a short duration, it transitions to the BOTH-ON state. This enables the receiver to deal with two cases: (a) the sender is using one radio whereas the receiver is out-of-sync and listening on the other radio, (b) the sender is out of range of the current radio but in range of the other radio. If no packet is received in the BOTH-ON state, it implies that the sender has dropped out of contact of both radios, therefore the receiver switches back to the IDLE state.

Summary of benefits: Having described the sender and receiver state machines, we now briefly describe the main benefits of our switching protocol.

- *Active mode efficiency:* During a block transfer, we minimize the amount of time for which both radios are turned on at the sender and receiver. This ensures that our system almost always consumes only as much energy as a single radio system.
- *Idle mode efficiency:* A key advantage of our protocol is that during idle times, only one receiver is operating in duty-cycled mode and the other is off. Thus, idle listening cost is minimized.
- *Low packet overhead:* All state transitions in our protocol are triggered by flags set in data packets. There are no additional control packets, hence our protocol has extremely low packet overhead.
- *Robustness:* Our protocol is robust to channel vagaries and different mobility patterns, and can recover from lost packets/acks, disconnections, and out-of-sync states.

V. ARTHROPOD IMPLEMENTATION

We have built a prototype multi-radio platform called *Arthropod* and have implemented the Q-Learning based adap-



(a) Hardware Prototype

(b) Software Architecture

Fig. 3. System Components: (a)Hardware prototype comprising the Tinynode and a CC2420 expansion board and (b)Unified Link Layer for the radios

tive link-layer and switching protocol. This section describes the hardware and software implementation of our system.

A. Dual-radio Hardware

Our *Arthropod* sensor platform consists of a low-power microcontroller and a pair of heterogeneous low-power radios. The current prototype employs a MSP430 microcontroller, a CC2420 radio, and an XE1205 radio. Rather than constructing such a platform from scratch, we employed an existing *TinyNode* sensor platform [6], which contains a MSP430 processor and the XE1205 radio, and augmented it with a custom-built daughterboard comprising the CC2420 radio. Figure 3(a) depicts the resulting prototype hardware of *Arthropod*.

The particular choice of the XE1205 and the CC2420 radios was governed by their complementary range-power characteristics (as described in Table I). While the peak range of the XE1205 is 2 kms for a bandwidth setting of 1.2 Kbps and +15dBm power level, we were unable to get reliable transmission on the XE1205 at this setting due to known calibration problems with the TinyNode’s XE1205 radio. Therefore, we use a data rate of 38.1 Kbps @+15dBm, at which setting the maximum range is 800m. In contrast, the CC2420 cannot transmit beyond 0dBm and thus has a maximum range of 80m.

B. Multi-radio Protocol Implementation

The software implementation for *Arthropod* is an adaptive link-layer that unifies the individual MAC layers for the two radios. We have implemented a unified radio interface as part of the TinyOS-2.x operating system for motes [20].

TinyOS-2.x Radio Drivers: The interfaces to the CC2420 and XE1205 radios allow for fine-grain control of many parameters including link layer acknowledgements, clear channel assessment, radio channel selection, data rate and transmit power. Each radio MAC layer supplies feedback required by our unified link layer: whether or not a packet was acknowledged, and the number of congestion backoffs experienced for the current transmission attempt.

Unified Link Layer: At an application level, our system provides a unified link layer that determines the radio interface currently best suited for communication (shown in Figure 3(b)). To the programmer, there is a single interface to send and receive packets to and from the node.

Environment	Mobility Pattern	example scenario
urban-indoor	continuous w/ obstructions	people in a building
urban-outdoor	continuous partial LOS	moving vehicle
urban-outdoor	nomadic	bus w/ stops
foliage	nomadic	bus w/ stops

TABLE II
BRIEF SUMMARY OF DATASETS COLLECTED

Duty-cycling: Our implementation of the switching protocol uses the low power listen (LPL) protocol for duty-cycling radios [14]. In this approach, the sender can wakeup the receiver in a completely asynchronous manner by sending a long preamble that is at least as long as the sleep cycle of the receiver. The sender uses such a long preamble to “wakeup” the receiver and initiate the block transfer.

While LPL is available as part of both the CC2420 and XE1205 radio stacks in TinyOS 2.0, we experienced several problems with the XE1205 LPL implementation. For example, the XE1205 interface would occasionally silently drop a packet transmission while reporting successfully acknowledged delivery to the receiver. To circumvent this problem, the IDLE state switches off the long-range XE1205 radio and duty-cycles the CC2420 radio. This is not ideal for reasons outlined in Section IV — wakeup using the long-range radio ensures greater contact duration between nodes. We believe that this is a software problem with the XE1205 radio stack, and are working on fixing the problem.

VI. EXPERIMENTAL EVALUATION

In this section, we present a detailed evaluation of the Q-learning based unified link layer using a combination of experiments using data traces, results from Q-Learning running on an *Arthropod* mote, as well as implementation benchmarks. Our evaluation has three parts. First, we evaluate the performance of the Q-learning link layer in adapting to a diverse set of mobility patterns. Second, using collected traces we evaluate how well the learning algorithm handles power control across the two radio interfaces. Finally, we present benchmarks from an implementation of the link layer for an *Arthropod* mote to demonstrate that the described Q-Learning algorithm is efficient and has low resource usage.

A. Datasets

To ensure repeatable experimentation of the link layer, we gathered datasets under different conditions using our hardware prototype. We obtained four types of datasets that are a good representation of mobility patterns found in mobile sensor network deployments. Table II contains a brief summary of all datasets collected.

In each experiment, we configure the long-range XE1205 radio to a data-rate to 38.1 kbps and power level 15 dbm, whereas the short-range CC2420 radio transmits at the default 250 kbps at 0dbm. The traces are obtained by having both the radios transmit packets back to back at the rate of about 2 pkts/second. For each packet, the number of backoffs and

retransmissions are logged in the local flash memory of the sender, and retrieved later.

1) *Traces showing continuous mobility:* We obtained datasets with continuous mobility to represent two practical sensor application scenarios: wearable sensors and vehicular sensor networks.

- **Indoors:** This trace was collected indoors within our department. The receiver is stationary while the sender moves up and down the length of the corridor of the building (120 meters) transmitting packets while moving at a normal walking speed. The nodes are obstructed from each other by numerous walls as they move apart. The up-and-down movement pattern is repeated five times.
- **Outdoors:** This trace was collected on a stretch of road outside our department. The receiving node was placed on a bus stop shelter; the sender approaches the receiver from 600 meters, initially obstructed by buildings and foliage. The node briefly pauses at the bus shelter and continues down the street another 300 meters disappearing behind foliage and buildings. The sender moves at a rate of roughly 9 meters/second.

2) *Traces showing nomadic mobility:* Nomadic mobility refers to the case where nodes move for some time, pause at a specific location for a while, and continue in the same pattern. This type of movement behavior is common amongst sensors monitoring people or a habitat sensor network. We obtained two such traces:

- **Urban Outdoor:** This trace was collected at our campus near some large HVAC buildings and parking lots. The sender starts in close proximity to the receiver and moves away at a normal walking speed. The sender pauses for a minute at locations 50m, 60m, 80m, and 70m away from the receiver, and finally returns to the receiver location. Line of sight is limited during this trace resulting in poor performance for the CC2420 radio.
- **Habitat Outdoor:** This trace was collected outdoors in a wooded rural area with significant foliage. The sender starts 100 meters from the receiver and approaches the sender at a slow walking speed. The receiver pauses for 2 minutes near the receiver, and then moves away at a slow walking speed to a location 100 meters away. At all locations, significant number of trees introduce signal attenuation. The XE105 has good connectivity for the entire experiment, while the CC2420 moves in and out of range.

B. Trace-driven Evaluation of Q-Learning

These experiments evaluate the performance of the Q-learning algorithm in a MATLAB simulation environment and its performance for the various mobility traces described above. To get an accurate measure of the performance of the Q-learning based link layer, we emulate the behavior of the sender and receiver state machine (Section III) given the sequence of packet losses observed in the traces. (Later, in Section VI-D we show that this emulation accurately corresponds to the performance of the real protocol in practice.)




Fig. 4. Energy consumed per successful packet for each dataset and strategy




Fig. 5. Percent Packets lost for the two radio interfaces and Q-Learning implementation

For all datasets, we used an identical set of Q-Learning parameters: $\alpha = 0.1$, $\gamma = 0.7$ and $\epsilon = 0.025$. These parameters were chosen since they seem to work well across a range of mobility datasets.

Q-Learning Performance: Figure 4 and Figure 5 summarize the energy per successful packet transmission and loss rates observed by our adaptive multi-radio link layer in comparison with using just one of the radios. In terms of energy consumption, the Q-Learning approach reduces energy consumption compared to the XE1205 radio by an average of 27% (the maximum reduction is 53.6% for the Outdoor Bus dataset), while incurring roughly 2-4% increased loss across the four cases. The slightly increased loss rate of Q-learning is caused by exploring an alternate interface periodically and transient losses caused while the algorithm is still learning. Similar energy gains of a maximum of 62.5% and an average of 44.6% are obtained over an approach that just uses the CC2420 radio but the improvements in loss rate are significantly higher (25%-60%). The results show that in all cases, the energy consumption of the adaptive multi-radio link layer is better than exclusively using either the CC2420 or XE1205 radio, while keeping the link loss rate to be close to that observed by the long-range XE1205 radio.

As can be seen, the worst case for the Q-Learning protocol is the outdoor nomadic trace where our benefits are only marginal in terms of energy. This is because connectivity using the CC2420 radio is highly sporadic and also very lossy (65% loss). Thus, our link layer is unable to take advantage of the CC2420 radio due to the high dynamics on it.

radio/power-level	% packets lost	energy consumed
XE1205@0dBm	4.24	.659mJ/Tx Success
XE1205@15dBm	0	.925mJ/Tx Success
CC2420@-25dBm	37.01	1.1mJ/Tx Success
CC2420@0dBm	35.45	1.2mJ/Tx Success
Q-Learning	3.53	.430mJ/Tx Success

TABLE III
STATISTICS FOR MULTIPLE RADIO / POWER CONTROL TRACE




Fig. 6. Time series trace for indoor continuous mobility pattern showing cumulative energy consumption for each radio / power-level as well as energy used for the Q-learning algorithm

C. Algorithm performance for power control across radios

A logical extension to the unified link layer is handling transmission power control in addition to radio interface selection. The CC2420 radio is capable of transmitting packets from -25dBm up to 0dBm, while the XE1205 can transmit from 0dBm to 15dBm. Increasing transmit power will provide longer range connectivity but at a higher energy cost; the optimum strategy will choose the minimum transmit power level on the most efficient radio without significantly increasing loss rate.

To evaluate power control across radio interfaces, we collected a packet trace similar to the *Indoor Continuous* described earlier. In addition to logging retransmissions and backoffs for the XE1205@15dBm and CC2420@0dBm, we log similar statistics for the XE1205 and CC2420@0dBm and -25dBm respectively. The number of states in the Q-Learning algorithm increases from 2 to 4; we maintain a Q-value for each radio/power level combination. To reduce exploration overhead, we only explore the radio/power combinations above and below the current setting. Logically, the next setting expected from a mobile node would be one higher if the distance between sender and receiver has increased and one lower if the distance has decreased. Such an approach would scale even if there were more power states being considered per radio.

Table III summarizes the results and compares the Q-learning approach to just using one of the two radios at one of the power levels. As can be seen, Q-learning is 54% better in terms of energy consumption per successful transmission than only using the XE1205 radio at 15dBm but has comparable loss rate. The energy benefits over using the CC2420 radio are 64%; the loss rate also reduces by an order of magnitude.

Switching Protocol State	% Time Spent
HIGH ON	10.7
LOW ON	78.1
BOTH ON	11.2

TABLE IV
TIME AND ENERGY SPENT DURING DIFFERENT RECEIVE STATES




Fig. 7. Energy spent per packet by the sender and receiver. Labels on sender bars indicate packet loss rates.

Figure 6 shows a time series plot of cumulative energy consumed for each packet sent for the five schemes. The staircase shape of the CC2420 radios cumulative energy consumption is caused by the mobility pattern of the sending node; the portions of the plot with a steep slope indicate that the radio is consuming an increased amount of energy per packet because it is moving out of range. Overall, the Q-Learning based adaptive algorithm sends roughly 40% and 10% of the packets using XE1205 at 0 dBm and 15 dBm; and 25% of the packets on the CC2520 at -25 dBm and 0 dBm. The results validate that Q-Learning is able to utilize each power state opportunistically.

D. Implementation Results

To validate our implementation and show the performance of the radio switching protocol, we collect a new dataset with the same mobility pattern as the *indoor continuous* dataset. For this experiment, a pair of sender and receiver nodes run the switching protocol and Q-Learning algorithm online. This study aims to measure the actual per packet energy costs incurred by the sender and receiver. In particular, the receiver can become out-of-sync with the sender, resulting in the receiver turning both radios on, or timing out, all of which costs energy and results in more packet losses. Finally, we also breakdown the % time spent by the receiver in different states of the receiver state machine.

The maximum data rate achievable by our software implementation for a pair of nodes transmitting continuously is 70kbps. However, while logging packet statistics to the external flash, the data rate reduces to 14kbps. This lower data rate is a result of the Tynode platform multiplexing the SPI bus between the XE1205 radio and the external flash memory. As a result of this issue, the sender and receiver are not continuously sending data which causes idle gaps to appear between packets. This forces the receiver to expend additional energy while waiting for packets to arrive. Since the idle time is an artifact of our evaluation, we ignore these periods when

presenting results.

To understand the energy efficiency of our protocol at the sender, Figure 7(a) compares the energy consumed by a single radio strategy to that of the dual radio implementation. The per packet energy consumption numbers presented for the CC2420 and XE1205 only cases are from Section VI-B. The results in Figure 7(a) show that our adaptive algorithm is 64% more efficient than a CC2420-only scheme, and 43% more efficient than an XE1205-only scheme verifying the gains found in simulation. These energy efficiency gains are achieved while maintaining a loss rate of 1.6% which is not substantially higher than the XE1205 loss rate of 0.6% and much lower than the 43.0% loss rate of the CC2420 radio. These results validate our simulation study and show that substantial sender-side energy gains are achievable by opportunistically using the CC2420 radio, while providing a loss rate comparable to that of the XE1205 radio.

Figure 7(b) shows the amount of energy consumed at the receiver as a result of the decisions made by the sender. The energy efficiency of the receiver will always fall somewhere between the efficiency of the XE1205 and CC2420 radios, depending on how often each is used. Bringing up both radio interfaces is an unavoidable result of the radio switching protocol and represents overhead beyond that of a single radio strategy. Our evaluation shows that the dual-radio protocol used 70% less energy than the XE1205, but 13% more energy than the CC2420. It is important to note that the receiver uses an order of magnitude less power than the sender, which means the sender-side gains overshadow the receiver-side losses.

Finally, Table IV gives a breakdown of the percentage of packets the receiver spends in each state of the switching protocol as well as the energy-efficiency associated with each state. The receiver spends 10.7% of time in the HIGH-ON state, 78.1% of time in the LOW-ON state and 11.2% of time in the BOTH-ON state. Ideally the radio switching protocol will only force the receiver into the BOTH-ON state while exploring or handing off between radios. Exploration accounts for 4% of this time, while the other 7.2% is caused by explicit handoffs and timeouts.

E. Microbenchmarks

In this section, we briefly discuss measurement-based latency and energy consumption microbenchmarks based on our implementation of the unified link layer. The energy/latency overhead imposed on the CPU by our multi-radio adaptation algorithm implementation on the *Arthropod* is highly efficient and consumes less than a hundredth of the energy/latency of the radios used. This shows that the overhead introduced by software can be compensated by larger performance gains achieved through intelligent radio selection. The amount of memory overhead of our implementation is 111 bytes, which is a very small portion of the available 10kB. A much larger portion of program memory is required, however, because two radio stacks need to be instantiated; supporting an additional radio stack requires an additional 12kB resulting in a total

usage of 29kB out of the available 48kB of program memory, although we believe that this can be optimized considerably.

F. Parameter Sensitivity

All the results that we have described use an identical set of Q-Learning parameters: $\alpha = 1$, $\gamma = 0.75$ and $\epsilon = 0.04$. We found that a larger α value is generally helpful in mobility traces due to the need for fast switching. As the traces become more and more nomadic in nature (*i.e.* as they involve more waiting and less movement), the optimal choice of α reduces a little. However, $0.9 \leq \alpha \leq 1$ seems to be ideal in almost all settings. The choice of ϵ impacts how fast we can switch but it also impacts the energy consumption. A high ϵ can lead to more exploration overhead but is more reactive. We found that exploration roughly every 10 seconds or so provides a good balance but this can be tuned depending on expected dynamics. Finally, we found that the results were not very sensitive to γ , and works best in the range $0.5 \leq \gamma \leq 0.85$.

VII. CONCLUSIONS

In conclusion, we have made two major contributions in this paper. First, we have outlined the design of a new multi-radio sensor platform, the *Arthropod*, that pairs two radios - CC2420 and XE1205 - that offer diversity in frequency, power and range. Second, we have presented the design of a novel Q-Learning driven adaptive link layer that provides the abstraction of a single radio to applications using the radios, and transparently switches between radios depending on which radio offers the most energy-efficient communication channel. Our experiments using a number of interference and distance datasets confirm that the system can provide highly effective adaptation to a wide range of dynamics. Finally, we showed that the learning algorithm can be easily implemented with limited memory and computational overhead on a mote-class sensor platform.

REFERENCES

- [1] Y. Agarwal, R. Chandra, A. Wolman, V. Bahl, K. Chin, and R. Gupta. Wireless wakeups revisited: Energy management for voip over wi-fi smartphones. In *Proceedings of Mobisys*, Puerto Rico, USA, June 2007.
- [2] N. Banerjee, M. D. Corner, and B. N. Levine. An Energy-Efficient Architecture for DTN Throwboxes. In *Proc. IEEE Infocom*, May 2007.
- [3] N. Banerjee, J. Sorber, M. D. Corner, S. Rollins, and D. Ganesan. Triage: A Power-Aware Software Architecture for Tiered Microservers. In *Proceedings of Mobisys 2007*, Puerto Rico, USA, June 2007.
- [4] J. Burgess, B. Gallagher, D. Jensen, and B. N. Levine. MaxProp: Routing for Vehicle-Based Disruption-Tolerant Networks. In *Proc. IEEE INFOCOM*, April 2006.
- [5] R. Draves, J. Padhye, and B. Zill. Routing in multi-radio, multi-hop wireless mesh networks. In *Proceedings of the 10th annual international conference on mobile computing and networking*, pages 114–128, Philadelphia, PA, September 2004.
- [6] H. Dubois-Ferriere, R. Meier, L. Fabre, and P. Metrailler. TinyNode: A Comprehensive Platform for Wireless Sensor Network Applications. In *Information Processing in Sensor Networks (IPSN 2006)*, 2006.
- [7] H.-Y. Hsieh and R. Sivakumar. A transport layer approach for achieving aggregate bandwidths on multi-homed mobile hosts. In *Proceedings of ACM Mobicom*, Atlanta, GA, September 2002.
- [8] P. Juang, H. Oki, Y. Wang, M. Martonosi, L. Peh, and D. Rubenstein. Energy-efficient computing for wildlife tracking: Design tradeoffs and early experiences with zebranet. In *ASPLOS*, San Jose, CA, October 2002.
- [9] H. Jun, M. H. Ammar, M. D. Corner, and E. Zegura. Hierarchical Power Management in Disruption Tolerant Networks with Traffic-Aware Optimization. In *Proc. ACM SIGCOMM Workshop on Challenged Networks (CHANTS)*, September 2006.
- [10] N. Mishra, K. Chebrolu, B. Raman, and A. Pathak. Wake-on-WLAN. In *Proc. Intl Conf on the World Wide Web (WWW)*, pages 761–769, 2006.
- [11] V. Navda, A. P. Subramanian, K. Dhanasekaran, A. Timm-Giel, and S. Das. Mobisteer: using steerable beam directional antenna for vehicular network access. In *MobiSys*, New York, NY, USA, 2007.
- [12] T. Pering, Y. Agarwal, R. Gupta, and R. Want. CoolSpots: Reducing the Power Consumption of Wireless Mobile Devices with Multiple Radio Interfaces. In *Proc. ACM MobiSys*, pages 220–232, June 2006.
- [13] T. Pering, V. Raghunathan, and R. Want. Exploiting Radio Hierarchies for Power-Efficient Wireless Device Discovery and Connection Setup. In *VLSI Design*, pages 774–779, 2005.
- [14] J. Polastre, J. Hill, and D. Culler. Versatile low power media access for wireless sensor networks. In *Proceedings of the Second ACM Conference on Embedded Networked Sensor Systems (SenSys)*, November 2004.
- [15] A. Rahmati and L. Zhong. Context-for-wireless: Context-sensitive energy-efficient wireless data transfer. In *Proceedings of Mobisys*, Puerto Rico, USA, June 2007.
- [16] P. Rodriguez, R. Chakravorty, J. Chesterfield, I. Pratt, and S. Banerjee. Mar: A commuter router infrastructure for the mobile internet. In *Proceedings of the Second International Conference on Mobile Systems, Applications, and Services*, Boston, MA, June 2004.
- [17] E. Shih, P. Bahl, and M. J. Sinclair. Wake on Wireless: An event driven energy saving strategy for battery operated devices. In *Proceedings of the Eighth ACM Conference on Mobile Computing and Networking*, Atlanta, GA, September 2002.
- [18] T. Stathopoulos, M. Lukac, D. McIntire, J. Heidemann, D. Estrin, and W. Kaiser. End-to-end routing for dual-radio sensor networks. In *Proceedings of the IEEE Infocom 2007*, Anchorage, Alaska, USA, May 2007.
- [19] R. S. Sutton and A. G. Barto. *Reinforcement Learning: An Introduction*. The MIT Press, 1998.
- [20] TinyOS Website. <http://www.tinyos.net/>.
- [21] <http://www.maxstream.com>. XTend OEM RF Module for Long-range Data Transfer.
- [22] 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 (INFOCOM)*, pages 1567–1576, New York, NY, USA, June 2002.