

Radio Diversity for Reliable Communication in WSNs

Branislav Kusy, Christian Richter, Wen Hu, Mikhail Afanasyev, Raja Jurdak
Michael Brünig, David Abbott, Cong Huynh, and Diethelm Ostry
CSIRO ICT Centre, Australia
{brano.kusy}@csiro.au

ABSTRACT

Deployment of wireless sensors in real world environments is often a frustrating experience. The quality of radio links is highly coupled to unpredictable physical environments, leading to intermittent connectivity and frequent outages. Because link qualities are not predictable prior to deployment, current deterministic solutions to unreliable links, such as increasing network density or transmission power, do not adequately address this issue.

We propose a new dual radio network architecture to improve communication reliability in wireless sensor networks. Specifically, we show that radio transceivers operating at dual widely spaced radio frequencies and through spatially separated antennas offer robust communication, high link diversity, and better interference mitigation. We show through experiments that radio diversity can significantly improve end-to-end delivery rates, network stability, and transmission costs at only a slight increase in energy cost over a single radio.

Categories and Subject Descriptors

C.2.1 [Computer Systems Organization]: Computer-Comm. Networks – Network Architecture and Design

General Terms

Design, Experimentation, Reliability

1. INTRODUCTION

Low power wireless mesh is the most common communication architecture for wireless sensor networks (WSNs) today. In mesh networks, data is transported via a sequence of links which form a route between the source and destination nodes. The quality of radio links can fluctuate over time [22] due to changes in the environment and radio interference. State-of-the-art data collection protocols [8] are therefore designed to react quickly to changes in network

connectivity and repair their routing state. However, unless multiple alternative routes exist in the network, connection losses are unavoidable. Even if alternative routes exist, the re-routed traffic along with any additional control traffic may cause network congestion and increased packet collisions that can lead to poor responsiveness and throughput. Moreover, link failures become more expensive to repair with increasing network size as the routing state may need to be propagated network-wide.

In this paper we focus on improving primary radio link reliability as the key to allowing the construction of large reliable and efficient WSNs spanning regions with heterogeneous and time-varying environmental conditions. There are however only a limited number of ways in which link quality can be improved. Selection of node locations to achieve a robust wireless mesh is a challenging problem. Acceptable node sites are constrained by the application requirements [4, 25] and by the physical environment [13, 27]. Even if robust locations of the nodes are selected during deployment time, the network propagation characteristics may change over time, caused by seasonal effects, plant growth, constructions, or moving objects. Consequently, over-provisioning for better network robustness is a common practice today. For example, adequate performance in urban forest deployments [17, 24] may require node spacings as small as one-twentieth to one-fiftieth of the maximum datasheet radio range. Many deployments hand-place nodes in carefully selected locations to ensure a line of sight between them. High gain directional antennas are also used to construct a reliable communication backbone [11]. These practices increase the deployment cost by either requiring additional hardware or time-consuming iterative deployment adjustment procedures.

If on the other hand the node design provides robustness to the changing conditions, not only does the network become more reliable, the initial setup is of lesser importance. The main objective of this paper is to propose a new network architecture based on this radio communication diversity. We build a new platform with two independent radios operating at well-separated frequencies and spatially separated antennas and demonstrate that their combination can protect against communication loss from multi-path fading and environmental interference. This approach of considering Multiple-Input/Multiple-Output (MIMO) techniques for improving reliability in sensor networks is evidence of closing the gap between the sensor networks and wireless communication communities.

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

IPSN'11, April 12–14, 2011, Chicago, Illinois.

Copyright 2011 ACM 978-1-4503-0512-9/11/04 ...\$10.00.

Using frequency diversity generally incurs additional cost: lower energy efficiency, loss of compatibility with existing hardware, and increased cost due to additional components. Our implementation of a dual low power listening radio stack uses at most 33% more energy than a single radio stack. We argue the additional reliability achieved is well worth the incurred cost. Furthermore, if we consider the cost of operating sensor nodes as a whole, the overall power consumption, including the sensors or actuators that are connected to the nodes, will further decrease the overhead of idle listening. Compatibility with different hardware platforms is desirable to enable heterogeneous deployments. To avoid loss of compatibility we present an implementation that is fully compliant with single radio band communication. This not only eliminates this potential disadvantage but also opens up the opportunity to use the proposed setup to integrate hardware platforms that operate on different bands. The aspect of cost and size increases would primarily affect minimalistic systems that are not the chief target of our investigation: our focus is on universal sensor node designs.

The design of our node hardware platform is driven by practical as well as theoretical considerations. We selected two IEEE 802.15.4 compliant radio chips, operating in the 900 MHz and 2400 MHz bands and designed antenna mounts to compromise between optimum antenna spacing and practicality of the design. We use communication systems theory to argue that the resulting radio diversity significantly improves the robustness of our sensor nodes against multi-path fading. Moreover, our sensor nodes are unlikely to suffer simultaneous interference or multi-path losses in both bands as there are no man-made devices at present which generate simultaneous correlated signals in these bands. Our software contributions were implemented in TinyOS [1]. We developed a dual-band network stack in compliance with the hardware abstraction architecture [10]. We hide the radio diversity details behind hardware independent layers which allows existing TinyOS applications and system services to be ported to our platform without modification. At the same time, our optimized versions of network protocols, such as the 4-bit link estimator [6] and the collection tree protocol (CTP) [8], can utilize the full capabilities of the dual radios.

We tested the platform in a number of experiments. Our main evaluation scenario includes 30 nodes deployed in a region with indoor office spaces, outdoor locations containing heavy industrial machinery, and a forest with dense foliage. The nodes were subject to considerable multi-path fading and a dynamic environment (e.g., sporadic rain that resulted in a wet foliage and large temperature changes). We evaluated performance of data collection applications operating in three different modes: using only 900 MHz radios, using only 2400 MHz radios, and operating simultaneously in both bands. Even though all these networks exhibited high performance when the experiments commenced, both single-band networks failed to deliver data reliably from a number of nodes for extended periods of time during the experiments. The dual-band network was much more robust, achieving above 98% delivery rate for 100% of nodes.

The rest of the paper is organized as follows. After presenting related work in Sec. 2, we discuss the theoretical background for radio diversity in Sec. 3. Guided by the theoretical results, we present our hardware and software design in Sec. 4 and 5. We evaluate radio diversity in Sec. 6 and offer our concluding remarks in Sec. 7.

2. RELATED WORK

The use of multiple radios in high-end wireless data communication systems is a common technique referred to as MIMO. Although the multiple radios are usually integrated in a single chip, the analog RF circuits are separated for each antenna. Each transceiver circuit consists of digital to analog converter, low pass filter, mixer and power amplifier for the transmit side, and band pass filter, low noise amplifier, mixer, analog to digital converter, and poly-phase filter on the receiver side. An example of recent wireless LAN standards utilizing MIMO for improved range and data rates is IEEE 802.11n. However, the energy consumption of the currently available equipment is usually considered to be unacceptable for battery powered wireless sensor networks.

Existing wireless sensor network research indicates the potential of using different network interfaces for wireless communication [3]. The best choice for an application is usually based on a comparison of the intrinsic shortcomings of the different standards and hardware implementations.

Analysis and implementation of a range-diverse multi-radio system are described in [9] and [19]. In particular, [9] utilizes two radios which differ in RF output power and energy consumed per bit. However, modern RF chips offer a broad range of data rates and RF output power settings. With the addition of an adjustable external power amplifier, a fine grained control of the range and data-rate can be achieved by varying these settings. This in turn can minimize the energy expended per bit at the required range, without the cost of a second radio receiver. [19] discusses the use of sophisticated policies to decide when to switch to a different interface.

[2] and [23] discuss layered architectures where multiple radios or network interfaces are used to create a control- or coordination-channel in order to invoke higher level operations such as data transfer or processing. Unfortunately this results into a potential single point of failure - the control, or primary interaction channel. Interference or severe fading on this channel can result in a total loss of connectivity.

The responsibilities and opportunities that multiple communication interfaces impose on routing layers are illustrated in [5]. A new metric is proposed to address the shortcomings of the ETX metric (expected transmissions) and is implemented within a new MAC layer.

While the approaches reported in the cited work represent significant advances, their implementation would need modifications to existing network protocols to operate efficiently. There is also no possibility of integrating them with current networks that do not include multiple-radio links.

[28] illustrates different kinds of heterogeneity including link heterogeneity and the optimal placement of the heterogeneous nodes. We believe that such deployments are still error prone and unreliable over longer time periods. Especially in industrial environments, radio interference sources may change position and obstructions to the line-of-sight between nodes are usually quite dynamic. This means that the optimum placement for a node cannot always be decided at deployment time.

Active IEEE 802.11 access points or devices can interfere with sensor network deployments due to the overlap in channels and the significantly higher 802.11 RF output power permitted by the energy availability. [18] investigates this issue in detail and illustrates a multi-channel MAC layer. Recent proposals related to IEEE 802.15.4e include sophis-

ticated MAC layers for channel hopping. However, interference mitigation with these techniques does not address multipath fading, signal obstructions or interference over a complete frequency band.

To improve the reliability of links, [12] and [15] employ antenna diversity with specific algorithms for their selection. While the use of multiple antennas and polarization diversity can mitigate effects caused by obstructions or local fading, interference on a specific frequency can still cause the links to fail completely.

Deployments that build on the use of directional antennas usually assume that the link characteristics stay constant. In sensor networks, this assumption is often invalid due to the drastically changing environment over time.

Compared with the existing solutions and concepts, we are able to integrate fundamental diversity concepts from other wireless systems while introducing a minimal energy overhead. Our approach is able to mitigate the effect of RF interference occurring anywhere in a frequency band by using standard link metrics. We can adapt to static and dynamic line-of-sight obstructions through spatial diversity of the two antennas and their initial placement in the environment. In addition, our protocol implementation is able to interoperate with standard IEEE 802.15.4 compliant sensor nodes running the collection tree protocol of TinyOS.

3. THEORETICAL BASIS OF DIVERSITY

We use results from mobile radio and wireless LAN communications to analyze the benefits expected from antenna and frequency diversity for our platform. We conclude that the diversity increases robustness against signal loss due to multi-path, environmental fading, and interference across a wide range of environments.

3.1 Multi-path propagation

Multi-path propagation occurs when a radio signal is reflected and the reflections of the signal are detected by the radio receiver. If the reflections arrive out of phase at the receiver then destructive interference will reduce the signal detected by the receiver. The signal distortion from the multi-path depends on the distribution of regions in the environment that scatter or reflect the radio signal. We analyze two scenarios to determine the largest volume of space and the broadest band of frequencies in which signal degradation may occur.

3.1.1 Spatially selective fading

Spatial fading models are optimized for specific wireless environments as they depend on the angle and time of arrival of the multiple radio signals. The many existing models differ in the distribution of sources of scattering that reflect the radio signal. We consider two models: the far scattering model used in mobile radio network design and the near scattering model used in mobile handset and wireless local area network design.

The far scattering model.

This model assumes scattering that occurs a long way from the transmitter and receiver[16] (see Fig 1) and is derived from double slit interference. The interference pattern may be thought of as the correlation between two spatially separated copies of a radio signal. The first null in the inter-

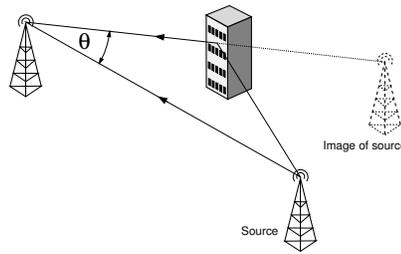


Figure 1: The far scattering scenario

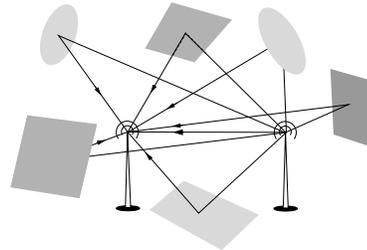


Figure 2: The near scattering scenario

ference pattern occurs, and radio signals are out of phase, when the distance separating the points is greater than:

$$d_{antenna} = \frac{1}{2} \frac{\lambda}{\sin \theta} \quad (1)$$

where λ is the wavelength of the radio carrier and θ is the spread in angle of arrival of the radio signals.

The near scattering model.

This model is valid when the radio signal is scattered close to and uniformly around the wireless device. The correlation $\rho_{electric}$ between radio signals separated by a distance of $d_{antenna}$ is given by the zeroth order Bessel function J_0 [21]

$$\rho_{electric} = J_0 \left(\frac{2\pi \times d_{antenna}}{\lambda} \right) \quad (2)$$

The first null in the Bessel function occurs at 2.5. That is

$$\rho_{electric} = 0 = J_0(2.5) = J_0 \left(\frac{2\pi \times d_{antenna}}{\lambda} \right) \quad (3)$$

$$d_{antenna} = \frac{2.5\lambda}{2\pi} \quad (4)$$

3.1.2 Frequency selective fading

Frequency selective fading is caused by differences in the times of arrival of multiple radio signals traveling along different paths. Radio signals will have independent multi-path fading if the frequency separation of the signals is larger than the coherence bandwidth of the environment, given by

$$B_{coherence} = \frac{2\pi}{\tau_{dispersion}} \quad (5)$$

where $\tau_{dispersion}$ is the channel dispersion, or arrival time difference between radio signals propagating along each path in the environment [16].

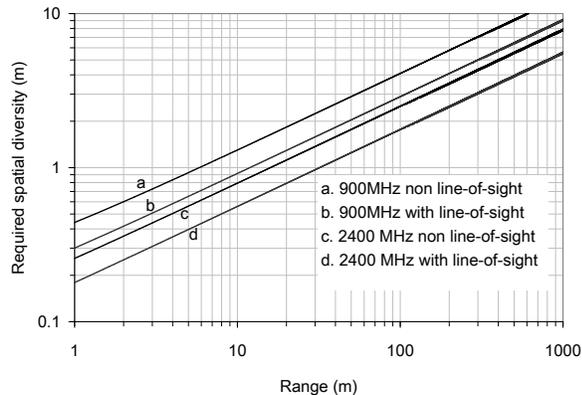


Figure 3: Calculation of maximum range at which antenna diversity provides benefit

3.1.3 Bounds on spatial and frequency fading

When the environment contains both near and far scattering volumes, the far scattering model provides a bound for the lowest rate of spatial fading. In Fig. 3, the lower bound for spatial fading is calculated. At larger path lengths the correlation between radio signals increases, thus the spatial separation between sampled signals must be increased.

Equation 5 shows that if the channel dispersion tends towards zero, for example, if scattering occurs from forest foliage close to one radio, the correlation between frequency separated radio signals increases and a wider frequency separation is needed to mitigate multipath fading.

3.2 Obstruction and environmental fading

In many deployments, sensor nodes are placed at or near the ground and the line of sight propagation path between the nodes is obstructed. Shadow fading occurs when approximately $\frac{1}{4}$ or greater of the first Fresnel zone of the radio signal is obstructed, the radii of the Fresnel zones given by

$$F_n = \sqrt{\frac{n\lambda d_1 d_2}{d_1 + d_2}} \quad (6)$$

where F_n is the radius of the n th Fresnel zone at a point located d_1 from one node and d_2 from the other node, and λ is the wavelength. Thus, the obstruction of the line of sight propagation path will be less at 2400 MHz than at 900 MHz

3.3 Interference

Interference is to be expected in the unlicensed spectrum that is used by many wireless devices. If the interference persists over longer timeframes or is periodic, the throughput of 802.15.4 radios decreases due to the limited access to the radio channel. Interference at 2400 MHz is mainly caused by WiFi networks. The WiFi channel occupies 20 MHz or 40 MHz of spectrum and many channels may be in simultaneous use. WiFi interference is able to saturate the entire spectrum available to the 802.15.4 radio [18].

In band interference at 900 MHz is not currently a significant problem for 802.15.4 radios in many regions. It is primarily caused by telemetry networks and cordless telephones. Telemetry signals occupy a narrow bandwidth which the direct sequence spread spectrum modulation is designed

to tolerate. Out of band interference at 900 MHz is caused by mobile telephones and pagers operating in adjacent frequency allocations.

3.4 Diversity design for multi-path

We explore the diversity requirements for wireless sensor nodes in three environments: (1) indoors in a semi open office space; (2) outdoors in open space; and (3) outdoors in a forested area with dense foliage. We apply the near and far scattering models to these environments and identify limitations of using only spatial or only frequency diversity in the sensor node design.

3.4.1 Limitations of spatial and frequency diversity

It is apparent from Figure 3 and Equation 5 that it is the scattering far from the radios that will determine the smallest antenna separation, while it is the scattering near to the radios that will determine the minimum frequency diversity that will increase reliability through diversity. Note that outdoor deployments in forested areas may be influenced by both the near and far scattering volumes.

Specifically, nodes are often placed close to objects that scatter radio signals in indoor deployments. If the scattering volume is located less than 1 m from a sensor node, the paths travelled by radio signals are at most 2 meters different in length, and Equation 5 dictates the frequency separation of at least 900 MHz to receive uncorrelated radio signals (802.15.4 spectrum usage is less than 85 MHz).

On the other hand, open-space outdoor deployments often cover long distances as there are few objects obstructing the direct line of sight. If the distance between nodes is 400 m and scattering occurs close to the midpoint, Figure 3 dictates the antenna spacing of at least 10 m to benefit from spatial diversity.

3.4.2 Using both spatial and frequency diversity

In our design, we exploit the tradeoff between the benefits of broadly spaced antennas and broadly spaced frequencies on one hand, and the cost penalty in size or energy when implemented in wireless sensor nodes. Our above analysis shows that antenna spacings of $\frac{\lambda}{4}$ that are suitable for indoor wireless deployments do not increase reliability in an outdoor deployment, while the spread-spectrum modulation of 2 MHz bandwidth in 802.15.4 (or even 20 MHz bandwidth in Wi-Fi) only provides limited benefit in the outdoor environment. If, however, the spatial and frequency diversity are used at the same time, the sensor node design may use a practical antenna separation of 1 meter and a frequency separation of at least 100 MHz to improve robustness of radio links against multi-path and frequency selective fading in all studied environments.

3.5 Diversity design for interference

Frequency diversity is robust against interference if the sources of interference are uncorrelated at the multiple frequencies. It is extremely likely that interference will be uncorrelated between the 900 MHz and 2400 MHz bands. No human radio sources other than our sensor networks currently transmit simultaneously at these two bands. This will not be true, if multi-band radio systems become popular. However, distributing network traffic across multiple bands will alleviate the interference in each band.

Table 1: RF chip specifications

Metric	AT86RF212	AT86RF230
Frequency	779-787MHz 863-870MHz 902-928MHz	2405-2480MHz
Data rates	BPSK: 20, 40 kbps, O-QPSK: 100, 200, 250, 500, 1000 kbps	O-QPSK: 250 kbps
TX power	-11 to 10 dBm	-17 to 3 dBm
Sensitivity (250kbps)	-101 dBm	-101 dBm
Link budget (250kbps)	111 dBm	104 dBm



Figure 4: A node mounted on a building wall

Assuming uncorrelated interference of the two signals, the probability of reception p of a signal simultaneously transmitted on dual bands is

$$p_{Dual} = 1 - (1 - p_{2400MHz}) \times (1 - p_{900MHz}) \quad (7)$$

For an example case of 90% reception of each signal the overall probability of error free reception would be 99%.

Alternatively it is possible to send packets over only one of the carriers using a MAC protocol to select the radio channel with the best characteristics. In this case the energy overhead comes only from the MAC protocol.

4. HARDWARE PLATFORM

We have developed a dual radio hardware platform that is based on our generic environmental sensing platform. We use Atmel Atmega 1281 low power MCU, the same MCU as used by the IRIS platform. The platform has a stackable design that allows us to easily connect and swap multiple expansion boards. We developed two radio boards: a 900 MHz board loaded with the Atmel AT86RF212 transceiver and a 2400 MHz board with AT86RF230 transceiver. Each board was configured to use a separate antenna and thus has its own RF connector. The features of the radio transceivers are shown in Table 1 for reference.



Figure 5: A node mounted in the forest environment

To protect the nodes against the environment, we used a waterproof housing. We mounted two antennas on a horizontal plastic tube that can be attached to a vertical pole with a T-connector (see Figures 4 and 5). The placement of antennas on the plastic tube was selected according to guidelines outlined in Section 3.1.1 to provide optimal spatial diversity. Cables of approximately 1 m length route the antennas to RF connectors on the sensor node housing. For effective radiation, we used antennas that provide a 5dBi gain on 2400 MHz and 2.2dBi gain on 900 MHz. We used low cost off-the-shelf antennas and the gain difference comes from the efficiency of the respective antenna designs.

5. IMPLEMENTATION

We integrated the support for multiple radios in the network stack of TinyOS 2.1 [1]. This effort includes developing a driver that supports simultaneous operation of multiple radios, a link estimator that tracks radio link qualities of neighboring nodes on multiple radios, and a data collection protocol that optimizes its performance using diverse radio links. The structure of our code is shown in Figure 6.

As we presented in Section 1, one of the design criteria for our platform was to fully comply with standard TinyOS networks operating in a single radio band. This allows our platform to fully integrate with existing 802.15.4 based sensor deployments, for example, to improve their robustness to environmental interference. The interoperability is achieved by our strict adherence to the standard TinyOS interfaces and packet header definitions in all network layers. We continue our discussion by providing implementation details of the main network stack software components.

5.1 Dual-band Radio Driver

In accordance with the hardware abstraction architecture of TinyOS [10], we chose to mask the underlying radio diversity to applications and operating system services. Unless a software component specifically requires to use a particular radio, the default radio is used. Consequently, existing applications work without changing a single line of code, while network protocols can optimize their performance through specialized interfaces.

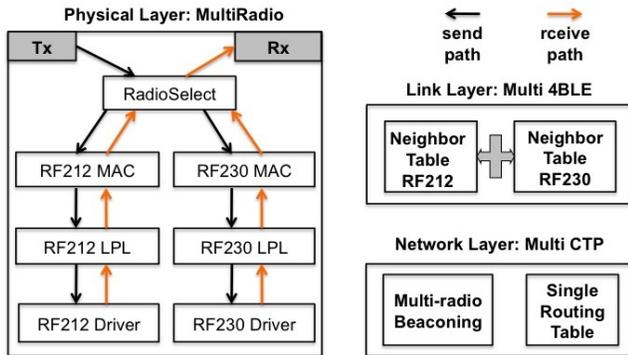


Figure 6: Software architecture of multi-radio network stack. The physical layer hides radio diversity behind standard interfaces. The link layer maintains a neighbor table per each band. The network layer makes routing decisions without explicitly considering radio diversity.

Drivers for both AT86RF212 and AT86RF230 radio chips are supported by TinyOS. Due to the similarity of the chip designs, the drivers are implemented to share most of their functionality through common software layers to save considerable system resources. Such tight coupling is undesirable if the two drivers are to operate simultaneously. For example, we need to be able to start and stop radios independently of each other to conserve energy during low power listening. We therefore separated the two radio stacks completely and executed the two drivers in parallel.

We further need to provide mechanisms to allow software components to either select a specific radio, or transparently choose the radio with the best chance of packet delivery for them. Our first principle is to store id of the selected radio within the radio message structure. Specifically, `radio_id` byte is stored in the metadata part of TinyOS messages ([14]), to avoid its transmission on the channel. We also provide a `RadioSelect` interface to be able to access and/or set the radio id for radio packets from any TinyOS component. Typically, application layer and lower level network protocols set the radio id for the transmitted packets. The radio driver accesses the radio id shortly before transmission to decide on which radio the packet will be transmitted.

5.2 Link Estimation Layer

The most common metric to measure a radio link quality is the number of expected transmissions (ETX)[7], defined as the inverse of the packet reception rate. Link estimators track bi-directional link ETX for all single-hop neighbors. ETX can be estimated from data traffic by keeping track of the acknowledged unicast packets. However, many times the data traffic is insufficient or is never routed through some neighboring nodes. Link estimators, therefore, periodically broadcast control beacons at each node with monotonically increasing sequence numbers. Gaps in the received sequences indicate the number of missed beacons.

We used 4 bit link estimator (4BLE) [6] and extended it to keep track of radio link qualities on multiple radio bands. The extended 4BLE-multi protocol maintains a separate neighbor table and sequence number for each radio band and parameterizes all internal functions as well as ex-

ternally exposed interfaces by `radio_id` identifier. Isolating radio bands from each other is important in situations where one band performs consistently better. Since the neighbor table is limited in size, links of the better band could dominate the table and slow down the recovery process in case a severe interference renders this band temporarily unusable.

Similarly to the multi-band radio driver design, 4BLE-multi hides the radio diversity details from higher level software components. In particular, its interfaces are fully compatible with the original 4BLE. If the application does not request the link estimate on a specific band, link estimator returns the best link-ETX from among all available network interfaces.

5.3 Data Collection Layer

Data collection protocols provide best-effort delivery of data to one of the sink nodes. We used the standard CTP protocol implementation of TinyOS [8]. CTP extends the single-hop link ETX provided by the link estimator to multiple hops, forming a set of implicit trees that are rooted at the sinks. CTP defines multi-hop ETX as follows: ETX of the root is 0 and ETX of any other node is the sum of the ETX of its parent and the ETX of the link to the parent (provided by the link estimator). Parent is a neighboring node with the lowest multi-hop ETX. Nodes forward data packets through their parents, thus minimizing the cost of radio transmission in terms of the ETX metric. The CTP implementation is very tightly coupled with 4BLE to save system resources. In particular, 4BLE intercepts beacon packets sent by CTP and extends them with link specific headers.

CTP-multi, our extension of CTP to multiple radio bands, requires minimal modifications as it is completely agnostic to which radio band was used to calculate the multi-hop ETX. Recall that the multi-hop ETX is defined as the sum of the ETXs of all links en-route towards the sink. For each individual link, the 4BLE-multi estimator selects the optimal radio band that minimizes the link's ETX to the parent. The beauty of this approach is that even though the radio band details are hidden from CTP, minimizing multi-hop ETX during the parent selection implicitly finds the best routing paths across all available radio bands.

One modification that CTP-multi requires is the support for sending beacons on multiple radio bands. Since we expect uncorrelated interference and environment dynamics across bands, beaconing intervals on the different bands may vary. The control traffic is driven by CTP and *not* the link estimator, thus the beaconing logic of CTP needs to be duplicated for the multiple bands.

5.4 Low Power Listening

Long term operation of sensor networks requires sensor nodes to carefully manage their available energy resources. Numerous low power listening (LPL) protocols were developed to duty cycle radios [20, 29]. The default low power MAC protocol in TinyOS reduces idle listening by shifting energy burden from receivers to senders. Receivers alternate between a long sleep period and a short wake up period to check radio activity on the radio channel. If the data traffic volume is low, the time spent in the sleep period dominates operation of receivers, dramatically reducing their idle listening time. Since the receivers are in the sleep mode most

of the time, senders have to transmit long preambles to ensure that the receivers are awake to receive the packet.

The current radio drivers in TinyOS transmit the same data packet continuously during the whole preamble period. Therefore, the neighboring nodes, which are not the intended receivers, can go back to sleep after receiving the data packet without having to receive the whole preamble. Transmitters, on the other hand, can only go back to sleep after receiving an acknowledgement from the receiver during LPL preamble. An important optimization is to keep the receiving node in a listening mode for a short period of time (i.e., 10ms), after a successful reception of a packet. This allows the transmitter to deliver data to the receiver at high packet rates.

We completely separated LPL layers of the different radios, so that they can be started and stopped independently. Consequently, the cost of packet transmissions and receptions in our platform is exactly the same as in single-band radio platforms. However, the cost of idle listening will double as both radios need to be ready to receive a packet. We show that depending on the packet transmission rate, the energy overhead of idle listening increases the overall energy cost of radio transmissions by 3% - 33%.

6. EVALUATION

We evaluate benefits and trade-offs of using antenna spacing and radio band diversity in sensor networks in a series of experiments. We first conduct controlled experiments to demonstrate that radio diversity improves reliability of radio links in different environments. We then show the performance of our dual band platform in a large scale data collection experiment in an environment consisting of office space, outdoor open space, and forest environment. Finally, we quantify the energy overhead of dual-band networks for different data rates and LPL intervals.

6.1 Characterization of Diversity

We present outdoor and indoor controlled experiments to confirm that antenna spacing and radio band diversity improve reliability of radio links. We also confirm a negative result from recent literature and show that radio channel diversity provides only limited benefits in real-world urban deployments.

6.1.1 Outdoor Open Space Experiment

In our first experiment, we evaluate benefits of *antenna and frequency* diversity in an outdoor line-of-sight deployment with no measurable in-band interference. We used two pairs of nodes operating in 900 MHz and 2.4GHz bands and set their transmission power to achieve the same range for both radios (approximately 900 m). We mounted one node 1.5 m above the ground and placed another node on top of a car, elevated about 0.4 m above the roof to prevent intersection of the car with the Fresnel zone of the radios. We drove the car along a straight road with a small hill in the middle, stopping every 10 meters for tens of seconds. Each pair of nodes was continuously exchanging radio packets at 10 packets per second using 250 kbps data rate. We recorded the packet reception rate (PRR) using sequence numbers and the transmission distance using GPS.

For a given transmission distance, we plot the difference between measured PRRs in the two bands, as well as the individual PRRs in Figure 7. Each value in the figure is

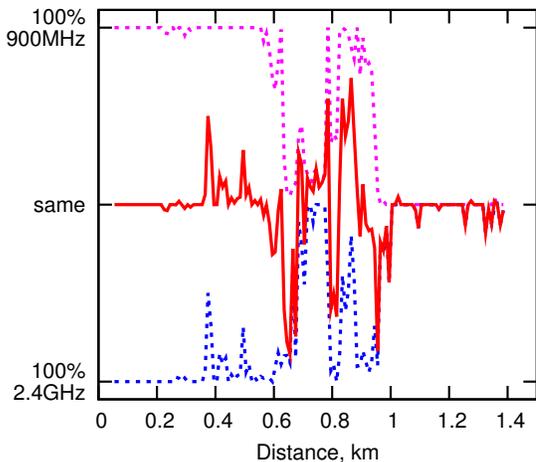


Figure 7: Difference between PRRs of 900 MHz and 2400 MHz bands over transmission ranges.

an average of multiple test runs measured along the same road. These results can be interpreted using the three-state characterization of wireless links [30]: the initial 40% of the distances correspond to the *connected* region with both radios attaining close to 100% PRR, the next 30% are in the *gray* region characterized by extreme variability in PRR due to small location change of the receiver, and the last 30% of the distances are in the *disconnected* region, with both radios achieving small PRR. We focus our attention on the gray region, where the benefits of radio diversity are maximized. Poorly performing radio links that belong to this region can be detected, flagged, and ignored by network protocols [30]. Our experiments in Figure 7 show that dual band radio can improve PRR of these links by as much as 84%. Fewer links then need to be ignored and the network connectedness improves significantly.

6.1.2 Indoor Office Space Experiment

Our second experiment evaluates the benefits of *frequency* diversity in an indoor environment. To eliminate the effect of spatial diversity due to two antennas, we used dual band antennas with the same 2.1 dBi gain on both frequency bands. We set the output power of the transmitter on both radios to 3dBm to ensure that the results are not distorted by the radios operating at the edge of their transmission range.

For both receiver and transmitter, the antenna connectors for the two bands are connected to a splitter and then to the dual band antenna. We measured the power loss of the cables, connectors and the splitter individually with a network analyzer and found, surprisingly, that the power losses were different in the different bands. After a more detailed study, we found that the dual band antenna was not very well matched to the transmitter in the 2400 MHz band, corresponding to a quite large power loss of 6 dBm. We used a fixed attenuator on the receiver side of 900 MHz to offset this measured difference and ran a few line-of-sight experiments to verify that equal power (RSSI) was received in both bands.

We placed a fixed transmitter in the center of an office environment. The transmitter sent ten packets per second on both radios. We sampled around 100 different locations and

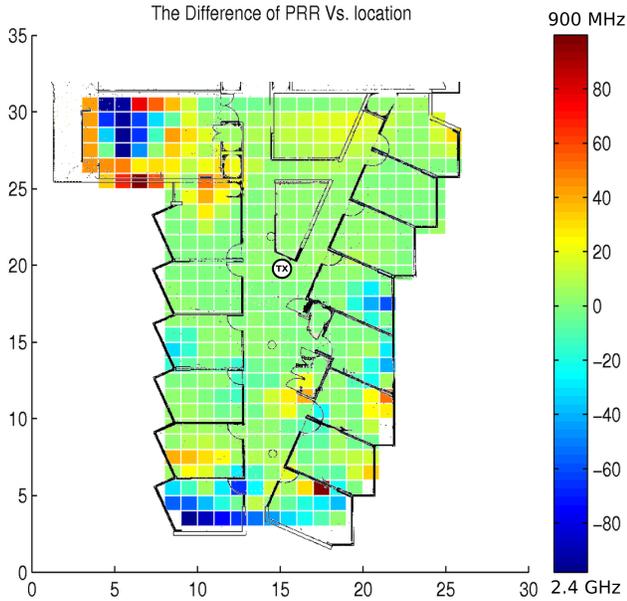


Figure 8: The PRR difference of 900 MHz and 2400 MHz bands, overlaid with the building floorplan. Red indicates 900 MHz performing better while blue shows 2.4 GHz outperforming 900 MHz.

measured PRR with a mobile receiver. All measurements are averaged over periods of more than 10 seconds each.

Figure 8 shows the results of the measurements. We observe a number of areas where one frequency band performs better, while the other band is better in a neighboring cell. In line-of-sight measurements, both bands perform similarly as can be seen in the center of the plot. The area with the most fluctuations was an electronic laboratory in the upper left corner. Since the figure shows relative performance of the two bands, it is impossible to see how well the radios performed in the absolute terms. We note that at least one of the radios performed close to 100% in the majority of these experiments.

6.1.3 Channel Diversity

The last of our controlled experiments aims to demonstrate that channel diversity is insufficient in a number of scenarios. Similar results have been recently presented in Ortiz et al [18]. This work refutes the general belief in usefulness of channel hopping schemes that lead to numerous channel hopping MAC protocols (e.g., [26]). We reach similar conclusions in observing that channel hopping provides only limited benefits when multipath fading affects the radio signals and is far less useful in avoiding external interference than the popular belief suggests.

Table 2: Fading correlation between two radio frequencies

Environment	Channel Frequencies (MHz)		
	918MHz 924MHz	924MHz 2450MHz	2435MHz 2450MHz
Indoors	0.88	-0.72	0.68
Forest	0.60	-0.16	0.22

Table 3: PRR with an interference source on f , g where $f = 2437MHz$ and $g = 906MHz$

$f+8MHz$		$f+23MHz$		$f+38MHz$	
μ	σ	μ	σ	μ	σ
0.10	0.25	0.80	0.01	0.98	0.01
$g+9MHz$		$g+12MHz$		$g+18MHz$	
0.80	0.01	0.93	0.01	0.99	0.01

We first compare the benefits of channel and band diversity by measuring the temporal correlation in the radio signal strength (RSSI) between closely spaced and widely separated radio signals. We use dual-band radios to collect results in indoor and dense forest environments in Table 2. Band diversity provides communication paths that show independent fading in the forest environment and anti-correlated fading in the indoor environment. Channel diversity, on the other hand, does not significantly increase the reliability of communications indoors, nor in the 900 MHz frequency band in the forest environment. The correlation of received signal strength fluctuations indicates the correlation in packet reception rate between frequency channels and bands.

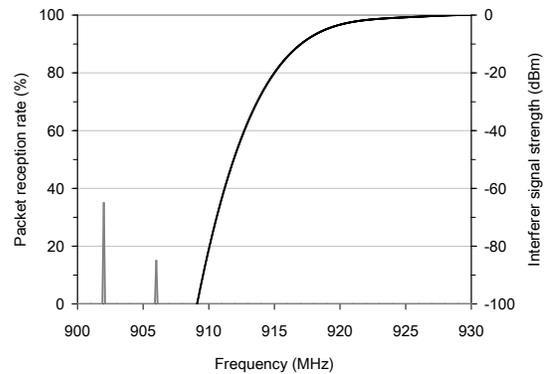


Figure 9: Measured packet reception rate (black) with GSM interference (grey)

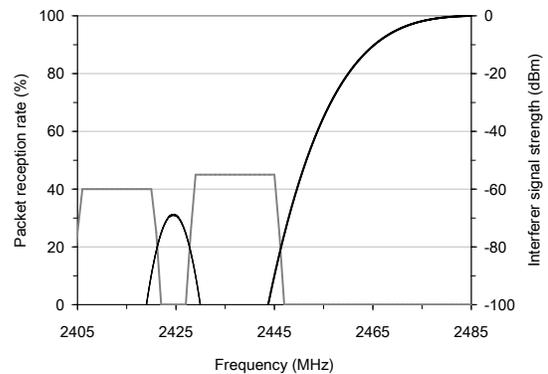


Figure 10: Measured packet reception rate (black) with WiFi interference (grey)

We also quantify the impact of the interference on nearby channels by measuring PRR of radio links in an indoor office space. The interference sources were a WiFi LAN transmitting at 2412 MHz and 2437 MHz and a GSM network operated at 906 MHz. As shown in Figure 9 and in Figure 10, high rates of packet loss occurred on both bands due to the MAC layer backoff in the IEEE 802.15.4 radios. We repeated the experiments with a single interferer only and provide quantitative evaluation in Table 3.

Specifically, the external interference caused 20% of the test packets to be dropped on both bands when a neighboring radio channel was used. When the sensor nodes operated with both radios enabled, the packet loss was reduced to the minimum packet loss of the two tested channels. In real world office environments, it is rare to see fewer than three non-overlapping WiFi channels active. Figure 10 suggests that no 2400 MHz 802.15.4 channel will provide sufficient robustness against interference in this case.

6.2 Data Collection Performance

We continue our evaluation of radio band and antenna spacing diversity in a complex real world environment over longer stretches of time. We compare a single-band 900 MHz network, a single band 2400 MHz network, and a dual-band network in a series of experiments. The placement of sensor nodes is identical for all three networks – we reprogram all nodes with the respective single- or dual-radio network stack before the experiments commence.

We use the hardware platform described in Section 4, with common off-the shelf antennas and both radios transmitting at 0dBm. The actual transmitted power at a given band depends on the transmission power of the radio, antenna gain, cabling, connectors, and how well these components are matched. Rather than carefully optimizing this part, we tried to reproduce a setup that a typical user might experience: using cheap off-the-shelf components and default radio settings that balance power consumption and radio range.

6.2.1 Experimental Setup

We measured the range of both radios with line of sight at approximately 900 meters for reliable communication. Equipped with these performance indicators, we deployed 30 nodes in a 400 m by 450 m area at our campus. It is a challenging environment (see Figure 11 for the deployment overview and Figures 4 and 5 for examples of node locations), containing 18 multi-story buildings equipped with wireless computers and access points, and cordless and mobile phones; industrial areas where heavy machinery is operated (such as hot metal carriers used in the mining industry); radar and satellite towers; and a forest with dense foliage, a creek, and elevation variations. We were precluded from deploying nodes in some areas due to safety constraints.

Given these challenges, we devised a deployment plan so that the maximum distance between neighbor nodes is 80 meters, one-tenth of the range achieved in unobstructed environments. All nodes are stationary during experiments, but are exposed to real-time link dynamics from surrounding environment as well as external interference.

Our test application periodically sends data packets to a base station using CTP. Instead of filling data payload with dummy bytes, we use data packets to carry performance statistics of nodes en-route to the base station. The originating node includes the following statistics in the packet:



Figure 11: Deployment of 30 nodes in a combined indoor/outdoor environment. The triangle illustrates the base station.

- (a) number of transmitted and received data packets,
- (b) number of transmitted and received beacons, and
- (c) the parent rate change (churn).

In addition, each forwarder adds statistics for each hop that the data packet traverses:

- (a) ID of the forwarding node,
- (b) ID of the used radio,
- (c) link-ETX value, and
- (d) LQI of the packet upon reception.

We set the packet size to 90 bytes and inter-packet interval to 5 seconds. Consequently, the aggregate traffic rate at the base station is 4 kbits per second which is well below the link saturation. We test three networks (900 MHz, 2400 MHz, and dual band) on three consecutive days, started at the same time of the day. The weather was consistent on all three days. The network is reset before every experiment and we ignore the first 20 minutes of data to let link neighbor tables and routing tables stabilize.

6.2.2 Metrics

We evaluate performance of the test application using standard metrics:

Reliability: measured as end-to-end delivery rate, or the percentage of packets delivered to the base station.

Path length (PL): routing path length calculated as the number of hops a packet takes.

Total cost: the total number of transmitted data packets.

Hop cost: the number of transmitted data packets per delivery (including retransmissions).

6.2.3 Comparison of Dual to Single Band Networks

Evaluation of the high level statistics is shown in Table 4. It is clear that both single band networks have problems coping with link dynamics. Even though we verified that

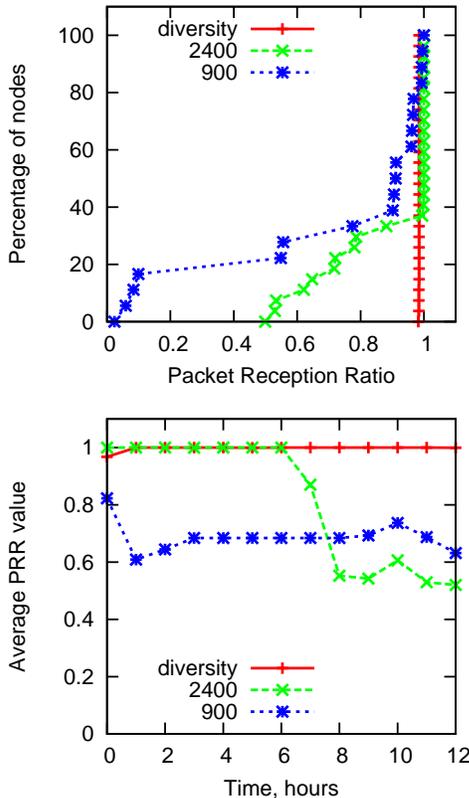


Figure 12: CDF of end-to-end delivery rates and its evolution over time. The three graphs correspond to three separate experiments.

all nodes were running and sending data during the first hour of the experiment, end-to-end delivery rates dropped significantly over time. The dual-band network, on the other hand, performed very well throughout the whole experiment. It is interesting to note that all networks used routes with similar path lengths. The increased hop cost in the single band cases indicates that even though the routes were of the same length, they required more packet transmissions to deliver packets to the base station. Note also that the PL metric is skewed in single-band cases, as we only average path lengths for the successfully delivered packets.

Next, we look at the high level statistics in more detail. Figure 12 shows the cumulative density function of end-to-end delivery rates for individual links and tracks average end-to-end delivery over time. We observe that the decrease in reliability of single-band networks is due to a few nodes that fail consistently. We verified in the data logs that these sets of nodes were different for different bands. In most

Table 4: Performance comparison of 900 MHz, 2400 MHz and dual band networks in the 30 node experiment.

	Reliability	PL	Total cost	Hop cost
Dual band	99.76%	2.24	736145	2.46
900MHz	70.39%	2.25	813151	5.61
2400MHz	87.27%	2.17	813429	3.52

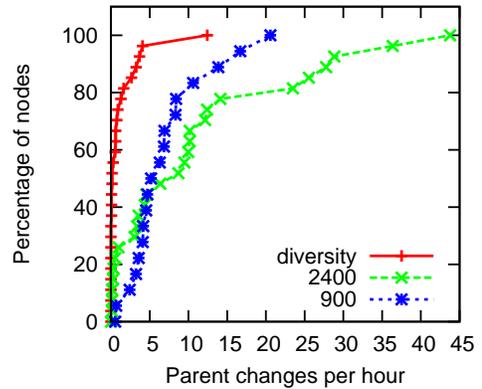


Figure 13: Stability of parent node selection. Higher congestion in the single band networks increased control traffic and decreased stability.

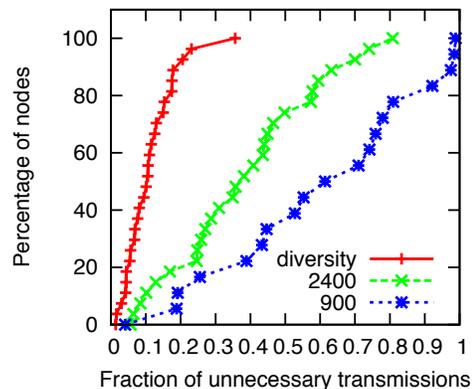


Figure 14: CDF of unnecessarily transmitted packets. Reasons for packet retransmissions include external interference and multipath fading.

cases, these nodes were parts of a network partition that became intermittently isolated during the experiment due to their limited and unstable link connections with the rest of the network.

We also study path lengths in more detail, as the overall averages of path lengths were somewhat surprising. We verified in our logs that some dual band routes are actually consistently slightly longer compared to their single band counterparts. However, we conjectured that even though these routes were shorter in terms of hops, they actually required more transmissions to reach the base station. This conjecture was confirmed by inspecting the aggregate link ETX for delivered packets where the dual band performed better than both single bands. Due to the lack of better options, the single band networks were forced to use these shorter, but less reliable routes.

In addition to the standard metrics, we study stability of routing paths and the efficiency of CTP in finding routes to the base station in Figures 13 and 14, respectively. CTP is designed to use stable routing paths – a new node becomes a parent only if its ETX is significantly better than the ETX of the old parent. Instability of routing paths, therefore,

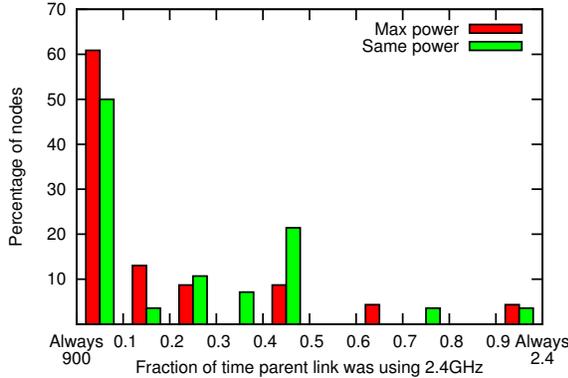


Figure 15: Evaluation of radio diversity. Dual band experiments with the same and the maximum transmission power. Eight percent of the nodes favor 2.4GHz radio even in the max power case.

signals high link dynamics to the point where considerable system resources need to be expended to repair the network. We can see in the figures that both the stability of routing paths and the fraction of packets never delivered to the base station are significantly worse in single-band networks.

6.2.4 Evaluation of Network Diversity

Finally, we evaluate the extent to which radio diversity is used in the dual band network. Recall that we do not enforce the diversity in software, it is solely up to the link layer to select better performing links. Transmission power of the two radios in the previous experiment was set to the same value, thus the lesser path loss for the 900 MHz band and the interference from WiFi on 2400 MHz band caused preference for the 900 MHz band. In fact, only 23.9% of transmissions were in the 2400 MHz band. We show in Figure 15 the fraction of nodes that used 2.4GHz radio to communicate with their parent.

To confirm the usefulness of diversity in a more extreme case, we repeated the above experiment using the maximum power settings on both chips. Due to the different maximum power limits of the radio chips, the relative transmit power of the 900 MHz radio increased by 7dBm. Even though the diversity in the second experiment decreased, some nodes still used 2.4GHz radio to a large extent. In fact, about 8% of nodes used 2.4GHz radio for majority of their transmissions.

6.3 Energy Overhead of Multiple Bands

In this section, we show that our platform does not incur a significant energy overhead when using two radios. Measuring energy consumption in outdoor large scale experiments is impractical, thus we focus our analysis on a controlled experiment in the lab. We use the default low power listening MAC layer (described in Section 5.4) and test a simple application with a single transmitter periodically sending unicast packets to a receiver. We measure the current drawn by the transmitter using a digital oscilloscope averaged over a 10 minute period. The sleep current of our platform was measured at $300\mu\text{A}$, while the active currents in the RX mode were 14.1mA and 19.1mA for RF212 and RF230, respectively. The active currents include the MCU active current

Table 5: Average energy overhead of dual radio stack relative to a single radio (RF230) under different data rates and LPL intervals.

Data Rate	Low Power Listening Period			
	0ms	128ms	512ms	1024ms
0pps	55.5%	33.3%	25.0%	20.0%
1pps	55.5%	27.8%	22.3%	18.0%
10pps	56.0%	14.5%	13.0%	12.3%
100pps	48.6%	3.3%	3.3%	3.3%

(approximately 4mA), thus the expected worst-case overhead of the two radios over RF230 is 10.1mA, or about 53%.

We conduct experiments for four different LPL periods and four different data rates and report our findings in Table 5. If the node keeps its radio continuously on (LPL=0ms), the overhead of the dual stack is as high as 56%. However, the overhead decreases to below 33% when we use LPL and transmit data. In particular, the energy overhead of the dual radio stack is minimal under high data rates as the power consumption is completely dominated by one radio transmitting data packets. In large scale deployments, nodes may need to forward packets for other nodes that are further away from the base station. Note that above calculations do not consider the energy required to receive packets. However, the additional time the nodes spend in the receive mode will improve the energy overheads in Table 5, as again, only one radio is active when receiving a packet.

We argue that this simple experiment provides a basis for fairly accurate estimation of the energy overhead of higher layer network protocols. In particular, CTP uses only unicast transmissions to deliver data packets to the base station and control traffic is a small fraction of the data traffic [8].

7. CONCLUSIONS

We have introduced a network architecture that aims to improve reliability of primary radio links by making use of radio diversity. We have developed a sensor network platform that provides frequency diversity through two 802.15.4-compatible radio transceivers operating in 900 MHz and 2400 MHz radio bands. Additionally, we mounted multiple antennas on each node to gain additional resilience against multipath fading in different environments. We show through theoretical modeling and controlled experiments that our platform can effectively mitigate radio signal fading due to interference as well as scattering volumes located both near and far away from the sensor nodes.

We implemented our diversity network stack and hardware drivers in TinyOS. One of the main features of our implementation is inter-operability of our software with existing TinyOS applications and system components. The diversity network stack is thus fully compliant with existing 802.15.4 TinyOS networks.

We evaluated our platform in a large scale deployment of 30 nodes. The nodes occupied a mixed indoor and outdoor environment and were exposed to varying weather elements. We demonstrated that the dual-radio stack achieves significant improvements in end-to-end delivery rates, data transmission costs, and delivery latencies over single-radio networks deployed in the same locations. Furthermore, the improved reliability comes at a moderate energy overhead of less than 33% using low power listening.

Our future work includes further improvement of the energy consumption of the diversity radio stack. By turning off the radios that are not being used, we hope to be able to further decrease the overhead of idle listening. The solution to this problem is not obvious, as turning off one of the radios compromises radio diversity in the first place. However, either the leaf nodes, or the nodes whose majority of children use the same band are obvious candidates for energy optimization. We would also like to utilize the different transmission and reception costs of the different radio chips and optimize the selection of the transmission radio for the energy cost, in addition to the ETX value.

8. REFERENCES

- [1] TinyOS 2.1.0. <http://www.tinyos.net/tinyos-2.1.0/>.
- [2] J. Ansari, X. Zhang, and P. Mähönen. Multi-radio medium access control protocol for wireless sensor networks. *SenSys '07*, Nov 2007.
- [3] P. Bahl, A. Adya, J. Padhye, and A. Walman. Reconsidering wireless systems with multiple radios. *SIGCOMM Computer Communication Review*, 34(5), Oct 2004.
- [4] C. Á. Conaire, K. Fogarty, C. Brennan, and N. O'Connor. User localization using visual sensing and RF signal strength. In *Proc of ImageSense'08*, 2008.
- [5] R. Draves, J. Padhye, and B. Zill. Routing in multi-radio, multi-hop wireless mesh networks. In *MobiCom '04*, pages 114–128, 2004.
- [6] R. Fonseca, O. Gnawali, K. Jamieson, and P. Levis. Four-bit wireless link estimation. In *Hotnets-VI*, 2007.
- [7] O. Gnawali. TEP 124: The Link Estimation Exchange Protocol. <http://www.tinyos.net/tinyos-2.x/doc/>.
- [8] O. Gnawali, R. Fonseca, K. Jamieson, D. Moss, and P. Levis. Collection tree protocol. In *SenSys '09*, pages 1–14, New York, NY, USA, 2009. ACM.
- [9] J. Gummesson, D. Ganesan, M. Corner, and P. Shenoy. An adaptive link layer for range diversity in multi-radio mobile sensor networks. *INFOCOM '09*, pages 154–162, 2009.
- [10] V. Handziski, J. Polastre, J. Hauer, C. Sharp, A. Wolisz, D. Culler, and D. Gay. TEP 2: Hardware Abstraction Architecture. <http://www.tinyos.net/tinyos-2.x/doc/>.
- [11] J. Hicks, J. Paek, S. Coe, R. Govindan, and D. Estrin. An easily deployable wireless imaging system. In *Proc of ImageSense'08*, 2008.
- [12] R. Ketcham and J. Frolik. A low-complexity, compact antenna for mitigating frequency-selective fading. *IPSN '07*, Apr 2007.
- [13] S. Kim, S. Pakzad, D. Culler, J. Demmel, G. Fenves, S. Glaser, and M. Turon. Health monitoring of civil infrastructures using wireless sensor networks. In *IPSN '07*, pages 254–263, New York, NY, USA, 2007. ACM.
- [14] P. Levis. TEP 111: message_t. <http://www.tinyos.net/tinyos-2.x/doc/>.
- [15] J. Liang and Q. Liang. Channel selection algorithms in virtual mimo sensor networks. *HeterSanet '08*, May 2008.
- [16] A. Mehrotra. *Cellular Radio Performance Engineering*. Artech House, Norwood, MA, USA, 1994.
- [17] R. Musaloiu-E, A. Terzis, K. Szlavecz, A. Szalay, J. Cogan, and J. Gray. Life under your feet: Wireless sensors in soil ecology. In *EmNets '06*, 2006.
- [18] J. Ortiz and D. Culler. Exploring diversity: evaluating the cost of frequency diversity in communication and routing. *SenSys '08*, Nov 2008.
- [19] T. Pering, Y. Agarwal, R. Gupta, and R. Want. Coolspots: reducing the power consumption of wireless mobile devices with multiple radio interfaces. *MobiSys '06*, Jun 2006.
- [20] J. Polastre, J. Hill, and D. Culler. Versatile low power media access for wireless sensor networks. In *SenSys '04*, pages 95–107, New York, NY, USA, 2004. ACM.
- [21] M. Schwartz, W. Bennet, and S. Stein. *Communications Systems and Techniques*. McGraw-Hill, New York, NY, USA, 1966.
- [22] K. Srinivasan, M. A. Kazandjieva, S. Agarwal, and P. Levis. The β -factor: measuring wireless link burstiness. In *SenSys '08*, pages 29–42, New York, NY, USA, 2008. ACM.
- [23] T. Stathopoulos, M. Lukac, D. McIntire, J. Heidemann, D. Estrin, and W. Kaiser. End-to-end routing for dual-radio sensor networks. *Proceedings of the IEEE Infocom*, 2007.
- [24] R. Szwedczyk, E. Osterweil, J. Polastre, M. Hamilton, A. Mainwaring, and D. Estrin. Habitat monitoring with sensor networks. *Commun. ACM*, 47(6):34–40, 2004.
- [25] T. Wark, W. Hu, P. Corke, J. Hodge, A. Keto, B. Mackey, G. Foley, P. Sikka, and M. Brunig. Springbrook: Challenges in developing a long-term rainforest wireless sensor network. In *Proc of ISSNIP'08*, 2008.
- [26] T. Watteyne, A. Mehta, and K. Pister. Reliability through frequency diversity: why channel hopping makes sense. In *PE-WASUN '09*, pages 116–123, New York, NY, USA, 2009. ACM.
- [27] N. Xu, S. Rangwala, K. K. Chintalapudi, D. Ganesan, A. Broad, R. Govindan, and D. Estrin. A wireless sensor network for structural monitoring. In *SenSys '04*, pages 13–24, New York, NY, USA, 2004. ACM.
- [28] M. Yarvis, A. Kushalnagar, H. Singh, Y. Liu, and S. Singh. Exploiting heterogeneity in sensor networks. In *Proc. of the IEEE Infocom*, 2005.
- [29] W. Ye, J. Heidemann, and D. Estrin. Medium access control with coordinated adaptive sleeping for wireless sensor networks. *IEEE/ACM Trans. Netw.*, 12(3):493–506, 2004.
- [30] C. Yin and A. Terzis. Poster abstract: On the spatial characteristics of the gray region for 802.15.4 radios. In *IPSN '09*, pages 393–394, Washington, DC, USA, 2009. IEEE Computer Society.