

# Empirical Study of a Medical Sensor Application in an Urban Emergency Department

JeongGil Ko  
Computer Science  
Department  
Johns Hopkins University  
jgko@cs.jhu.edu

Tia Gao  
AID Networks  
tia@aidn.org

Andreas Terzis  
Computer Science  
Department  
Johns Hopkins University  
terzis@cs.jhu.edu

## ABSTRACT

User needs and technology availability drive the introduction of wireless sensing applications in clinical environments. While these applications have the potential to improve efficiency and quality of care, very little is known about their performance during day-to-day use at the hospital. In this work, we use data from a deployment of a 802.15.4-based wireless sensor network at the Emergency Room of the Johns Hopkins hospital to answer these questions. Specifically, over a period of ten days we deployed a system of wireless vital signs monitors that measure the heart rate and blood oxygen levels of Emergency Room patients. During this time we collected statistics about the network's RF links, the performance of its tree routing protocol, and its end-to-end reliability. We find that the hospital environment we tested has considerably higher radio noise levels across multiple frequency channels and more bursty links compared to other indoor environments. Nonetheless, the routing protocol we use finds high quality links and the end-to-end packet reception ratio is above 99.9%. Taken as a whole, these preliminary results suggest that despite the challenges that clinical environments pose, wireless medical sensing applications can perform well in these conditions.

## 1. INTRODUCTION

The US healthcare system faces a variety of growing challenges, including an aging population, increasing chronic disease prevalence, nursing shortages [1], and increased concern with response to natural and man-made disasters. All these challenges require a hospital system that is better able to provide effective care to large numbers of patients under the best scenarios and to potentially respond to a surge of patients in emergencies. These user needs, coupled with the recent availability of technology that can automate some of the repetitive tasks performed manually today, are driving the introduction of wireless sensing applications in healthcare. However, the lack of results from deployments in clinical settings is creating a vicious cycle in which healthcare

providers are reluctant to offer their patient care areas for experiments with unproven technologies leading to technologists not being able to test their systems in realistic conditions.

In this work we attempt to resolve this conundrum through results collected from a pilot deployment of a medical sensing application deployed at the Emergency Room (ER) of the Johns Hopkins hospital. Specifically, with the help of our partners at the Emergency Department of the Johns Hopkins School of Medicine, we deployed *MEDiSN*, a system for monitoring the vital signs of ambulatory patients. *MEDiSN* comprises multiple *miTags*, which are custom-built motes that periodically measure the patients' heart rates and blood oxygen levels then transmit them using their IEEE 802.15.4 radios [21]. A wireless backbone, also part of *MEDiSN*, delivers these measurements to one or more gateways. *MEDiSN* was deployed over a period of ten days in November/December 2008 throughout the waiting area of the Johns Hopkins ER to monitor the vital signs of patients waiting to be seen by the doctors.

Based on the data from this deployment, we summarize our contributions as follows. (1) We evaluate the radio channel conditions in a real clinical setting. Specifically, we measure the temporal characteristics of point-to-point 802.15.4 links over two different frequency channels. The results indicate that the ER shows elevated levels of channel noise that can cause bursty losses. Additionally we show that multiple sources of interference exist in the hospital environment. (2) We present the performance of a medical sensing application (*MEDiSN*) in an urban ER environment. We find that the tree routing protocol that we use can select good quality links and achieve end-to-end delivery ratios that are higher than 99.9%.

The rest of the paper is structured as follows. We outline the design of the *MEDiSN* system in Section 2 where we also present the application deployed at the Emergency Department of the Johns Hopkins hospital. Section 3 presents all the results from this deployment. We review related work in Section 4 and close in Section 5 with a summary.

## 2. MEDISN

Healthcare in the United States is facing a double threat. From one side, demand for healthcare is expected to increase as suggested by demographic trends for an aging population and increasing prevalence of chronic diseases. At the same time, supply is decreasing due to nursing staff shortages [1] and decreasing hospital capacities [22]. This combination leaves US healthcare systems facing immense challenges on

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.

*BodyNets '09* Los Angeles, California USA  
Copyright 2008 ICST 978-963-9799-41-7.



**Figure 1: Medical Information Tag (miTag):** A Tmote mini based physiological monitor that measures heart rates and blood oxygen levels. Vital signs are displayed on the on-board LCD screen and transmitted over the MEDiSN wireless backbone.

a daily basis. Moreover, there is growing concern about the hospitals' ability to provide effective care to a rapid influx of patients in emergency situations. Tools that automate and improve the patient monitoring process are helpful in any circumstances, but are particularly so in "surge" situations.

With these needs in mind, we developed *MEDiSN*, a wireless sensor network for automating the process of patient monitoring in hospitals [8]. *MEDiSN* consists of multiple *Physiological Monitors (PMs)*, which are battery-powered motes equipped with sensors for measuring patients' physiological data. *PMs* temporarily store collected measurements and transmit them after encrypting and signing them. Figure 1 shows a monitor (called *miTag*) that we developed for measuring patients' heart rates and blood oxygen levels. This *miTag* uses a version of the Telos mote [17], with the same MSP430 MCU and CC2420 802.15.4 radio. The *miTag* uses a Nellcor OEM PulseOx controller [16] to collect the patient's vital signs and a  $128 \times 160$ , 65K color LCD screen to display them.

Unlike previous systems (e.g., CodeBlue [13]) in which *PMs* also relay data, *MEDiSN* incorporates distinct *Relay Points (RPs)* which self organize into a bidirectional wireless tree connecting the *PMs* to one or more gateways. The division of functionality between acquiring and relaying data, enables *PMs* to achieve consistent, predictable behavior and low energy consumption, while allowing us to engineer the system to provide superior end-to-end service. Specifically, while *PMs* can be mobile, the *RPs* have pre-determined, fixed positions that allow us to provision a high-quality wireless backbone. Moreover, because *PMs* are not responsible for relaying traffic, they can aggressively duty cycle their radios to reduce energy consumption. On the other hand, *RPs* are powered through the hospital's electrical outlets and therefore do not have energy constraints. Figure 2 provides

an outline of the *MEDiSN* architecture. The software running on the *miTag* and *MEDiSN*'s *RPs* uses TinyOS 2.x [11].

The *RPs* form a routing tree, using the Collection Tree Protocol (CTP) provided by TinyOS [4], to forward the *PMs*' measurements to the gateway. CTP provides best-effort service but uses hop-by-hop retransmissions to increase the end-to-end packet reception ratio (PRR). *PMs*, on the other hand, do not use CTP. Doing so would potentially involve *PMs* acting as relays, countering the principles of the *MEDiSN* design. Instead, a *PM* sends its data to the *RP* that shares the best link with that *PM*. Furthermore, the *PM* retransmits a packet for a maximum number of times or until it receives an acknowledgment.<sup>1</sup>

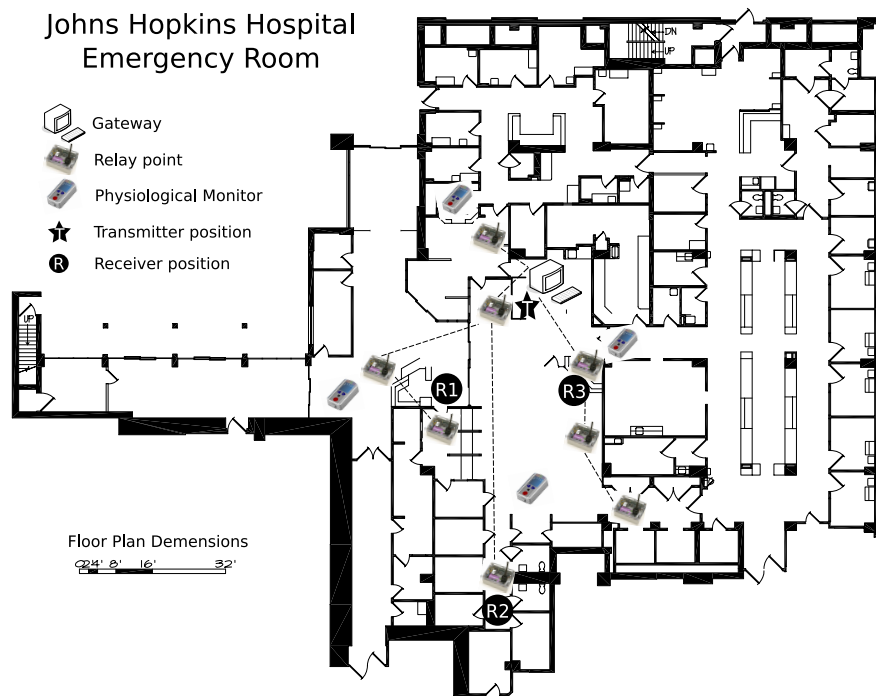
*RPs* also route messages from the gateway to individual *PMs*. To do so, each *RP* generates a Patient Information Packet (PIP) every five seconds and forwards it to the gateway. The PIP includes all the *PMs* which connect directly to that *RP*, as well as the *RP*'s identity. *RPs* that receive a PIP append their ID before forwarding it toward the gateway. Then, when the PIP eventually arrives at the gateway, it contains a list of *RPs* that can be used to reach a set of *PMs*. When the gateway needs to send a downstream message to one of these *PMs*, it generates a source route by reversing the list of *RPs* contained in the PIP message. Such downstream messages use the same hop-by-hop retransmission mechanism to increase the probability of the gateway's messages to the *PMs*. Moreover, since *PMs* duty cycle their radios, the last *RP* buffers the downstream message until the *PM* sends its next set of physiological data.

The PIP messages also include information about the quality of the end-to-end routing paths. Specifically, each *RP* includes the signal strength and Link Quality Indicator (LQI) values of its link with the *RP* it receives the PIP from. The included signal strength and LQI values correspond to the exponentially-weighted moving average (EWMA) of all the received packets from that link ( $\alpha = 0.5$ ). Signal strength is measured over the first eight symbols and reported as the Received Signal Strength Indicator (RSSI) in dBm. Signal quality (LQI) is also measured by the radio over the first eight symbols and is reported as a 7-bit unsigned value that can be viewed as the average correlation value or chip error rate [21]. In Section 3.2 we use these measurements to evaluate the quality of the *MEDiSN* end-to-end paths.

## 2.1 Emergency Room Monitoring

Overcrowding occurs in 40% of all Emergency Rooms (ERs) in the U.S., where patients wait on the average of 3.5 hours before being seen by a doctor. Although the cases are rarely publicized, there have been a number of deaths in waiting rooms of urban hospital ERs [6]. In cooperation with our partners at the Johns Hopkins School of Medicine, we performed a pilot deployment of *MEDiSN* in the Emergency Room of the Johns Hopkins Hospital. The goal of this Institutional Review Board (IRB)-approved deployment was to monitor the vital signs of patients who are scattered in the waiting rooms and hallways of the ER waiting to be admitted. During this study we repeatedly deployed *MEDiSN* between 6 p.m. and midnight for ten consecutive days, culminating with a continuous 24-hour deployment.

<sup>1</sup>Performance of CTP is out of the scope of this work. Details and extensive evaluation on its performance can be found in [4, 8]




**Figure 2:** The MEDiSN network comprises multiple Physiological Monitors (PMs) and a wireless backbone of Relay Points (RP) that provides bidirectional communication between the PMs and the gateway. The figure also shows the floor plan of the Johns Hopkins Hospital Emergency Department over which MEDiSN was deployed.

A typical hospital deployment starts with a brief site survey during which we select the locations of the RPs. Specifically, we first deploy the gateway which is a PC/laptop with a Telos mote connected to its serial port. Once the gateway is active it starts transmitting beacons that RPs use to join the network. We then incrementally deploy RPs in a way that increases the coverage area while ensuring that they can join the routing tree. We use visual cues (i.e., LEDs) to determine when an RP joins and disconnects from the network. Once all the RPs are deployed we briefly walk through the area with a miTag to ensure that no coverage holes exist. The whole process requires less than 30 minutes for an area of the size shown in Figure 2. Using the routing link statistics described above we can iteratively change the topology of the wireless backbone, by removing RPs that are not used (i.e., do not relay traffic) and adding RPs in areas that have poor connectivity (i.e. low LQIs).

Once the network is fully deployed, we start deploying PMs to patients located in the Emergency Room waiting area. Patients wear the device around their neck with the help of a lanyard, while a disposable clip similar to the one shown in Figure 1 is externally attached to their finger. The patients wear the device until they are admitted or until they leave the ER. During that time, patients can move freely throughout the waiting room area while their vital signs are continuously monitored.

### 3. EVALUATION

Next, we present experimental results obtained from the MEDiSN deployment at the Johns Hopkins Hospital Emer-



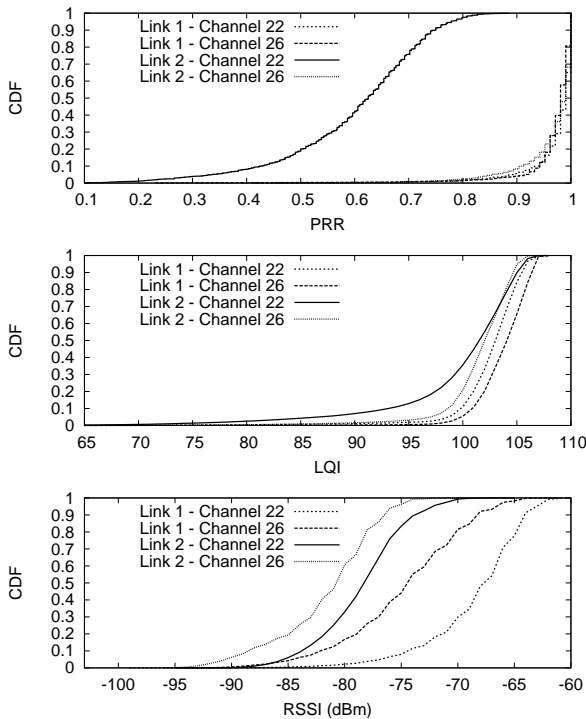
**Figure 3:** Overlap between WiFi and ZigBee channels in the 2.4 GHz frequency range.

gency Department in Baltimore, MD. As shown in the floor plan of the emergency room in Figure 2, the size of the ER waiting area where MEDiSN was deployed is approximately 70 feet × 90 feet.

#### 3.1 Channel Characteristics

The conditions of the wireless medium affect the performance of MEDiSN, as every other wireless medical sensing application. Given that MEDiSN uses ZigBee radios operating in the 2.4 GHz ISM band, we measure channel conditions in that frequency range. As Figure 3 illustrates, except for channels 25 and 26, all ZigBee channels overlap with one or more WiFi channels<sup>2</sup>. This observation is important because the Emergency Room at Johns Hopkins (as many large hos-

<sup>2</sup>In fact, in Europe and Asia where as many as 14 WiFi channels are available, *all* ZigBee channels overlap with WiFi.



**Figure 4: CDFs of per-minute PRR, LQI and RSSI values for  $L1$  and  $L2$ . PRR and LQI levels are higher for channel 26, while packets received on channel 22 have higher RSSI values.**


pitals) is covered by a hospital-wide WiFi network active on 802.11 channels 1, 6, and 11. Thereby, in order to investigate the impact of WiFi interference we perform experiments on ZigBee channels 22 and 26.

Specifically, for each channel we place one transmitter and three receivers at the same locations within the waiting area (shown in Figure 2 as  $T$  and  $R1, R2, R3$  respectively). The selected locations correspond to the locations of the gateway and some of the relay points (RPs) used by the MEDiSN network. In what follows,  $L_n$  is the link between the transmitter and receiver  $n$  ( $n = 1, 2, 3$ ). Because we use TelosB motes ([17]) for RPs, all eight motes are also TelosB motes.

The transmitter broadcasts<sup>3</sup> 111-byte packets (equal to the packet size used by MEDiSN) with an Inter-Packet-Interval (IPI) of 500 msec. We ran the test continuously over a period of 24 hours to explore whether channel conditions in the ER vary throughout the day. Each receiver records the RSSI and LQI values of every received packet. The RSSI is the RF signal strength in dBm, while the LQI is measured over the first eight symbols and is reported as a 7-bit unsigned value that can be viewed as the average chip error rate. Furthermore, the receiver collects an RSSI sample after receiving each packet. These samples are used to record the channel noise level.

Figure 4 presents the CDFs of the per-minute packet reception ratios (PRRs), average LQI and RSSI for  $L1$  and

<sup>3</sup>We note that senders do not perform any CSMA checks prior to broadcasting packets.



**Figure 5: Noise floor and packet RSSI values for  $L2$  on channel 22 and channel 26 over a 24-hour period. The noise level on channel 22 is  $\sim 10$  dBm higher than channel 26.**


$L2$ .<sup>4</sup> One notices that the PRR and LQI values for the same link are higher on channel 26 than 22. This result matches our intuition, as we expect channel 22 to be more “noisy”. Subsequent results will show this to be true. We would also like to point out that the difference in LQI values is less pronounced than the PRR differences. The reason is that LQI values are recorded only for successfully received packets and are therefore biased.

On the other hand, RSSI values for channel 22 are higher than channel 26 in Figure 4. Furthermore, the median RSSI value for  $L2$  on channel 22 is  $-78$  dBm, while the median PRR is only  $\sim 62\%$ . This result contradicts the findings of Srinivasan and Levis who found that links with RSSI  $> -87$  dBm have PRR  $> 85\%$  [20]. However, both observations can be explained through Figure 5, which plots the one-minute averages of the channel noise level and RSSI values. It is clear that the high RSSI values on channel 22 (see Fig. 4) are due to the elevated channel noise levels seen in Figure 5.


Figure 6 presents the per-minute PRRs and average LQI values for  $L2$  on channels 22 and 26. We make two observations based on this figure. First, channel conditions do not change appreciably throughout the day. This behavior is different from those observed in other indoor experiments (e.g., [7, 12]) and can be attributed to the fact that the ER is operational 24 hours a day. Moreover, we find that high LQI values are a necessary but not sufficient condition for high PRRs. While the LQI values for both channels 22 and 26 are comparable, the PRR varies significantly. Figure 7 that presents the correlation between PRR and LQI values for the two channels clearly illustrates this difference. This result is somewhat surprising<sup>5</sup> and contradicts previous studies which showed that links with high LQI values

<sup>4</sup>The performance of  $L1$  and  $L3$  were very similar so only  $L1$  and  $L2$  are presented throughout the section.

<sup>5</sup>We performed the same experiment with multiple receivers to ensure that the results were not due to the decreased hardware sensitivity of a particular receiver. All results showed the same behavior.



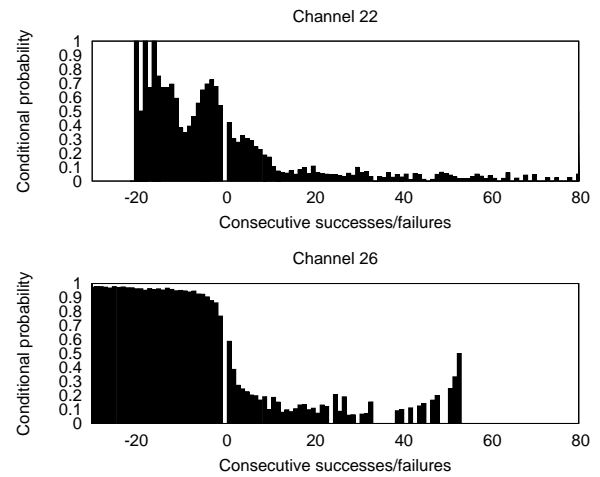
**Figure 6: Per-minute PRR and average LQI values for  $L_2$  over a 24-hour period. High LQI values are a necessary but not sufficient condition for low packet loss rates.**



**Figure 7: Correlation between per-minute PRR and average LQI values for  $L_2$  over a 24-hour period. While the range of LQI values observed in channel 22 is comparable to that in channel 26, the PRRs in channel 22 are considerably lower.**

have high PRRs [12, 20]. On the other hand these studies were performed on channel 26 and match the results from channel 26 in Figures 6 and 7. We believe that the PRR on channel 22 is lower, despite the high LQI values, due to bursty WiFi interference. This conjecture is supported by the high fidelity noise level measurements presented later in this section.

The measurements so far have shown that links can have intermediate PRRs over time. However, because medical sensing applications require high delivery ratios, RPs attempt to retransmit failed packets (i.e., packets that do not generate an acknowledgment). Next, we investigate the effectiveness of this strategy by performing the same test as before but with  $IPI = 50$  msec. Figure 8 shows the conditional packet delivery function (CPDF) generated from this experiment. The CPDF (introduced in [9]) corresponds to the probability of a packet being successfully received after  $n$  consecutive failures or successes. Consecutive successes are expressed as negative numbers, while consecutive failures are expressed as positive numbers. For example,  $CPDF_a(5)$  is the probability of a successful delivery after five consecutive



**Figure 8: Conditional Packet Delivery Function (CPDF) for  $L_2$  and  $IPI = 50$  msec. Failures on both channels exhibit a bursty behavior.**

losses on link  $a$ . Likewise,  $CPDF_a(-5)$  is the probability of a successful reception after five consecutive successes. In this way the CPDF provides a concise way of describing a link's burstiness. One can see from the non-uniformity of the CPDF plots in Figure 8 that compared to other environments (e.g., [10, 19]) the links at the ER are more bursty on both channels.

Summarizing the results so far, we found that links can be lossy and bursty, especially on channel 22. Furthermore, in a seeming contradiction to previous findings, lossy links can also have high LQI values. In order to understand the root cause of this behavior we measured the channel noise at a high fidelity. Considering that the transmission time of a maximum length 802.15.4 packet is  $\sim 3.5$  msec, we collect channel noise samples at a rate of 1 KHz. Figure 9 presents the results obtained over five minutes for both channels at the ER. The same figure also shows results from channel 26 collected at a university library in the night during final exams, a research lab environment in the evening, and a coffee shop during the day time. As expected, channel 22 at the ER shows higher average noise levels. More importantly the high frequency sampling process exposed noise spikes not seen in Figure 5. We conjecture that these spikes are due to interference from other radios (e.g., 802.11) and are the source of the links' burstiness.

We also note that, while not as pronounced as on channel 22, noise spikes can be seen in the ER trace from channel 26. Such spikes were largely absent from all the other environments that we tested. To further determine the nature of this noise, we performed a Fourier transform to the noise data collected from the ER and the coffee shop. Figure 10 presents the resulting power spectra. One can see that noise at the ER includes multiple periodic components, possibly from external interferers. On the other hand, the noise seen at the coffee shop can be modeled as white noise. While we do not know the source(s) of the noise seen in Figures 9 and 10 we believe that it is generated by other wireless devices operating in the 2.4 GHz ISM band, including cordless phones and Bluetooth devices.

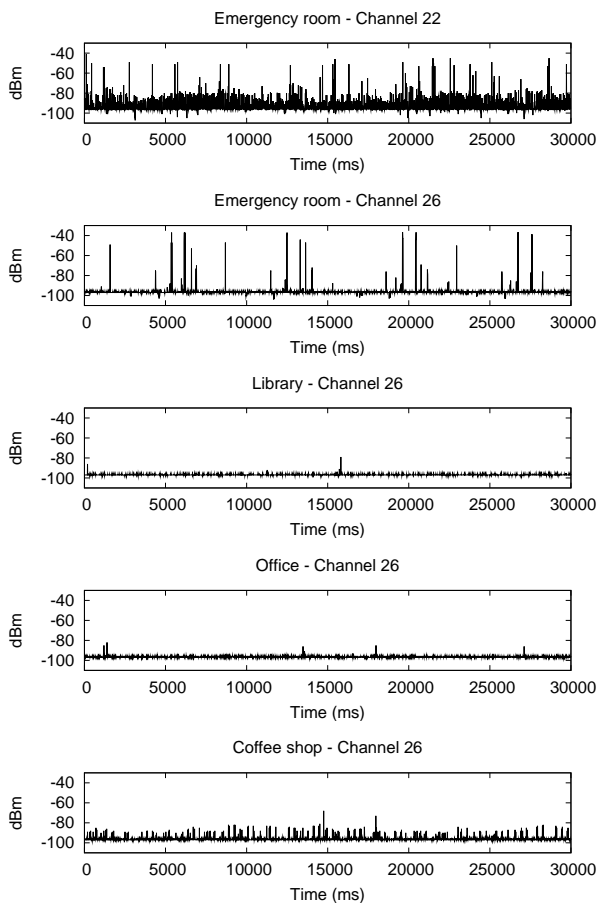


Figure 9: 1 KHz channel noise samples collected from both channels at the ER, a university library, research office, and coffee shop for channel 26. Both ER channels exhibit larger noise spikes than any other environment.

### 3.2 Routing Protocol Performance

Considering that quality (i.e., PRR) across links can vary significantly, it is interesting to investigate the quality of the links selected by CTP, MEDiSN’s underlying routing protocol for the wireless backbone. To do so, we collect statistics about these links using the collection mechanism described in Section 2. Statistics were collected by running two copies of the MEDiSN system simultaneously, one on channel 22 and another on channel 26, over a 24-hour period. The gateways and RPs of both systems were placed at identical locations, shown in Figure 2. Both systems carried vital sign measurements generated by miTags worn by ambulatory patients in the ER waiting room. In order to exercise the system further we injected additional traffic using TelosB nodes that mimicked the traffic that miTags generate. Each source generates one 111-byte packet every 500 msec and an average of three sources (PMs) on different locations were concurrently active on each channel.

Table 1 summarizes the characteristics of the links that CTP selected. The average, median and standard deviation were computed over all the links that existed at least once

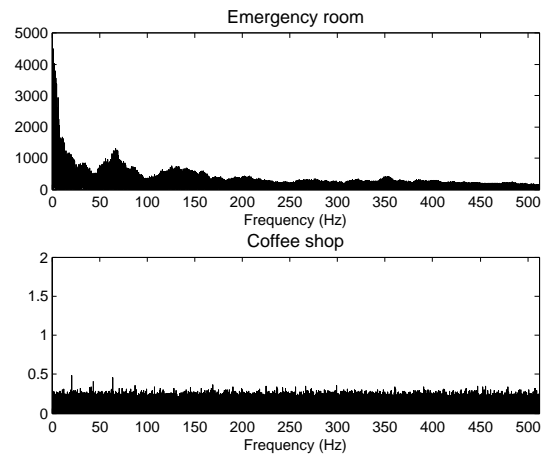



Figure 10: Power spectra of the channel noise measurements from channel 26 at the ER (top) and an urban coffee shop (bottom). The ER waveform has multiple strong frequency components, significantly deviating from white noise.

	Channel	Average	Median	Std. Dev.
LQI	22	98.74	101	9.81
	26	100.48	103	6.70
RSSI	22	-76.44	-78	13.35
	26	-75.42	-75	10.69

Table 1: Aggregate statistics for the routing tree links that MEDiSN uses at least once during a 24-hour test at the ER. CTP selects links with high RSSI and LQI values for both channels.

throughout the 24-hour testing period. Based on the results from Section 3.1, one can say that CTP selects high quality links at least for channel 26. While high LQI and RSSI values do not necessarily guarantee high PRR for channel 22, the application level results from the next section indicate that this is indeed the case.

CTP uses the 4-bit link estimator [5], which combines a variant of the ETX link metric [3] with information from the physical and data link layers. Furthermore, each time a node’s transmission fails (i.e., the node does not receive an acknowledgement), the node asks CTP for an alternate parent in the routing tree. If such an alternative exists, the old link is removed and a new link is established. Based on this observation, measuring the lifetimes of the network’s links provides an indication of both the quality of the links that CTP selects as well the “effort” that it has to make to ensure high delivery ratios despite link errors. Figure 11 presents the CDF of the links’ lifetime for all the links used at least once during the 24-hour test. We observe that the majority of the links have short lifetimes, reflecting changes due to lost packets. Furthermore, links on channel 22 have shorter lifetimes compared to links in channel 26. This discrepancy can be explained by the bursty behavior of links in channel 22 described in the previous section.



**Figure 11: CDF of the lifetime of all the links that existed at least once throughout the 24-hour experiment. Most links exist for short periods of time. Shorter lifetimes on channel 22 reflect the reaction of MEDiSN to the more adverse link conditions.**

### 3.3 Application Level Performance

During the 24-hour experiment described above, each PM transmitted 174,577 packets on the average. Of these, an average of 157.33 packets were lost for channel 22 and 42.67 for channel 26. These losses correspond to average packet reception ratios of 99.91% and 99.98% respectively, with standard deviations of 0.034 and 0.002. Furthermore, each packet successfully received at the gateway required an average of 1.18 per hop transmissions on channel 22 and an average of 1.08 per hop transmissions on channel 26. The average hop count for channel 22 was 3.11 hops, compared to 2.71 hops for channel 26.

While the loss rate and number of retransmissions are slightly higher on channel 22 these application-level results are very encouraging, especially considering the existence of intermediate quality links shown in Section 3.1. These results can be attributed to two key factors. First, MEDiSN quickly discovers and reacts to link level losses. Second, because the deployment is relatively dense, once a loss is detected CTP can provide an alternate path that can deliver the packet to the gateway.

## 4. RELATED WORK

A number of studies have investigated the properties of low-power wireless links in indoor deployments. The pioneering study of Zhao and Govindan [24], used Mica2 motes [14] to investigate the spatial and temporal properties of low-power wireless links inside an office building. However, that study used motes that do not operate in the 2.4 GHz frequency range that MEDiSN uses. Cerpa et al. performed a similar study with Mica2 motes and used the collected measurements to propose a model for the temporal properties of low-power wireless links [2]. The goal of this study is not to model link properties but rather to evaluate the impact of harsh hospital environments on medical sensing applications.

Srinivan et al. measured the characteristics of 802.15.4 links [18]. They found that links are bimodal over smaller timescales, while packet reception ratios (PRR) span a larger range over longer periods. Moreover, they observed that 802.15.4 links with the same long term PRR can have drastically different short term behavior, an observation that led to the introduction of the  $\beta$ -factor, a metric that represents the burstiness of radio links [19]. The results in Section 3.1 suggests that 802.15.4 links in clinical environments are more bursty than links measured in other indoor environments.

Lee et al. observed that different environments exhibit a wide variation in the levels of interference from 802.11 networks and proposed a mechanism to model this interference [9]. Musăloiu-E. and Terzis proposed distributed techniques that a WSN using 802.15.4 radios can use to detect interference from overlapping 802.11 radios [15]. The underlying network protocols used in this study are not frequency-agile, but we show that interference from 802.11 networks and other RF sources can be a severe source of problems for medical sensing networks.

Finally, Woo et al. illustrated the impact of link metrics on the performance of tree routing protocols [23]. The CTP routing protocol ([4]) that MEDiSN uses, relies on the 4-bit link estimation algorithm which has been shown to select high-quality end-to-end paths [5].

## 5. SUMMARY

This work presents results from a pilot deployment of the MEDiSN vital signs monitoring system at the Emergency Room of the Johns Hopkins hospital. We investigate the temporal characteristics of point-to-point 802.15.4 links in the ER over 802.15.4 channels 22 and 26. We find that links in channel 22 are more lossy, while the LQI values of packets received in that channel are not significantly lower than those in channel 26. On the other hand, packets received over channel 22 have higher RSSI values. These elevated RSSI values are due to increased levels of channel noise, that are also the cause of bursty losses in channel 22. Furthermore, we find that, unlike other environments we tested, channel 26 in the ER shows increased levels of non-random noise. We conjecture that other devices operating in the 2.4 GHz ISM band, such as cordless phones and Bluetooth devices, generate this interference.

We also collect measurements from the routing tree used to deliver patient data. We find that CTP, the tree routing protocol that MEDiSN uses, can find high quality links in both channels, but links tend to be short-lived. Nonetheless, the end-to-end delivery ratio is  $> 99.9\%$  for both channels and the average number of per-hop retransmissions is less than 0.2. Taken as a whole, these preliminary results are encouraging, as they show that despite the unique challenges that the emergency room environment poses to wireless sensing applications, they can provide high levels of performance.

## Acknowledgments

We would like to thank Yin Chen, Răzvan Musăloiu-E. and Jong Hyun Lim for their invaluable help in collecting some of the RF measurements. We are also grateful to Alex Ke-cojevic and all the staff at the Johns Hopkins Emergency Medicine Department for their support and to Dr. Richard Rothman for leading the ED pilot study for MEDiSN. Finally, we thank the anonymous reviewers for their insightful comments. This work was supported by the U.S. Department of Homeland Security (Grant Number N00014-D6-1-0991) through a grant awarded to the Center for Study of Preparedness and Critical Event Response (PACER) at the Johns Hopkins University. Any opinions, findings, conclusions or recommendations expressed in this publication are those of the author(s) and do not represent the policy or position of the Department of Homeland Security.

## 6. REFERENCES

- [1] American Association of Colleges of Nursing. Nursing Shortage Fact Sheet. Available at: <http://www.aacn.nche.edu/Media/Backgrounders/shortagefacts.htm>, March 2004.
- [2] Alberto Cerpa, Jennifer L. Wong, Miodrag Potkonjak, and Deborah Estrin. Temporal Properties of Low Power Wireless Links: Modeling and Implication on Multi-Hop Routing. Technical Report CENS-TR-44, University of California, Los Angeles, Center for Embedded Networked Computing, January 2005.
- [3] Douglas S. J. De Couto, Daniel Aguayo, John Bicket, and Robert Morris. A High-Throughput Path Metric for Multi-Hop Wireless Routing. In *Proceedings of the 9<sup>th</sup> ACM International Conference on Mobile Computing and Networking (MobiCom 2003)*, September 2003.
- [4] R. Fonseca, O. Gnawali, K. Jamieson, S. Kim, P. Levis, and A. Woo. The Collection Tree Protocol (CTP). TEP123 Draft. Available at: [http://tinysos.cvs.sourceforge.net/\\*checkout\\*/tinysos/tinysos-2.x/doc/html/tep123.html](http://tinysos.cvs.sourceforge.net/*checkout*/tinysos/tinysos-2.x/doc/html/tep123.html), April 2006.
- [5] Rodrigo Fonseca, Omprakash Gnawali, Kyle Jamieson, and Philip Levis. Four-Bit Wireless Link Estimation. In *Proceedings of the sixth workshop on Hot Topics in Networks (HotNets)*, November 2007.
- [6] Laura Gore. Emergency Patients Have Died Because of Crowded Conditions. American College of Emergency Physicians. Available at: <http://www3.acep.org/pressroom.aspx?id=33258>, October 2007.
- [7] Gregory Hackman, Otav Chipara, and Chenyang Lu. Robust Topology Control for Wireless Sensor Networks. In *Proceedings of the ACM Sensys*, November 2008.
- [8] JeongGil Ko, Răzvan Musăloiu-E., Jong Hyun Lim, Yin Chen, Andreas Terzis, Tia Gao, Walt Destler, and Leo Selavo. Demo Abstract: MEDiSN: Medical Emergency Detection in Sensor Networks. *Proceedings of ACM Sensys 2008*, November 2008.
- [9] HyungJune Lee, Alberto Cerpa, and Philip Levis. Improving Wireless Simulation Through Noise Modeling. In *Proceedings of the Sixth International Conference on Information Processing in Wireless Sensor Networks (IPSN'07)*, 2007.
- [10] HyungJune Lee, Alberto Cerpa, and Philip Levis. Improving wireless simulation through noise modeling. In *IPSN '07: Proceedings of the 6<sup>th</sup> IPSN Conference*, pages 21–30, New York, NY, USA, 2007. ACM.
- [11] Philip Levis, David Gay, Vlado Handziski, Jan-Heinrich Hauer, Ben Greenstein, Martin Turon, Jonathan Hui, Kevin Klues, Robert Szewczyk, Cory Sharp, Joseph Polastre, Philip Buonadonna, Lama Nachman, Gilman Tolle, David Culler, and Adam Wolisz. T2: A Second Generation OS For Embedded Sensor Networks. Technical Report TKN-05-007, Telecommunication Networks Group, Technische Universität Berlin, 2005.
- [12] Shan Lin, Jingbin Zhang, Gang Zhou, Lin Gu, John A. Stankovic, and Tian He. ATPC: Adaptive Transmission Power Control for Wireless Sensor Networks. In *Proceedings of the 4<sup>th</sup> ACM Sensys Conference*, 2006.
- [13] David Malan, Thaddeus Fulford-Jones, Matt Welsh, and Steve Moulton. CodeBlue: An Ad Hoc Sensor Network Infrastructure for Emergency Medical Care. In *Proceedings of the MobiSys 2004 Workshop on Applications of Mobile Embedded Systems (WAMES 2004)*, June 2004.
- [14] MICA Motes. <http://www.xbow.com/>.
- [15] Răzvan Musăloiu-E. and Andreas Terzis. Minimizing the Effect of WiFi Interference in 802.15.4 Wireless Sensor Networks. *International Journal of Sensor Networks (IJSNet)*, 2008.
- [16] Nellcor Puritan Bennet Inc. OxiMax NELL-1: OEM Pulse Oximetry Module. Available at: <http://www.nellcor.com/Serv/manuals.aspx?ID=291>, 2006.
- [17] Joseph Polastre, Robert Szewczyk, and David Culler. Telos: Enabling Ultra-Low Power Wireless Research. In *Proceedings of the Fourth International Conference on Information Processing in Sensor Networks: Special track on Platform Tools and Design Methods for Network Embedded Sensors (IPSN/SPOTS)*, April 2005.
- [18] K. Srinivasan, P. Dutta, A. Tavakoli, and P. Levis. Some implications of low power wireless to IP networking. In *Proceedings of the Fifth Workshop on Hot Topics in Networks (HotNets-V)*, November 2006.
- [19] Kannan Srinivasan, Maria Kazandijeva, Saatvik Agarwal, and Philip Levis. The  $\beta$ -factor: Measuring Wireless Link Burstiness. In *Proceedings of the 6th ACM Conference on Embedded Networked Sensor Systems (SenSys)*, 2008.
- [20] Kannan Srinivasan and Philip Levis. RSSI is Under Appreciated. In *Proceedings of the 3<sup>rd</sup> Workshop on Embedded Networked Sensors (EmNets)*, May 2006.
- [21] Texas Instruments. 2.4 GHz IEEE 802.15.4 / ZigBee-ready RF Transceiver. Available at <http://www.ti.com/lit/gpn/cc2420>, 2006.
- [22] The Center for Disease Control. Press Release: Visits to U.S. Emergency Departments at All-Time High; Number of Departments Shrinking. Available at: <http://www.cdc.gov/od/oc/media/pressrel/r050526.htm>, May 2005.
- [23] A. Woo, T. Tong, and D. Culler. Taming the underlying challenges in reliable multihop wireless sensor networks. In *Proceedings of 1<sup>st</sup> ACM Sensys Conference*, 2003.
- [24] Jerry Zhao and R. Govindan. Understanding Packet Delivery Performance In Dense Wireless Sensor Networks. In *Proceedings of the ACM Sensys*, November 2003.