Summer 8-1-2013

# Difficulties and Opportunities in Building Resilient Clinical Monitoring Systems with Wireless Sensor Networks

Rahav Dor
*Washington University in St. Louis*

WASHINGTON UNIVERSITY IN ST. LOUIS

School of Engineering and Applied Science

Department of Computer Science and Engineering

Thesis Examination Committee:
Dr. Chenyang Lu, Chair
Dr. Roger Chamberlain
Dr. Raj Jain
Dr. Tom Bailey

# Difficulties and Opportunities in Building Resilient Clinical Monitoring Systems with Wireless Sensor Networks

by

Rahav Dor

A thesis presented to the School of Engineering
of Washington University in partial fulfillment of the
requirements for the degree of
Master of Science

August 2013

Saint Louis, Missouri

# Contents

# List of Figures

# List of Tables

# Acknowledgments

First and foremost I would like to acknowledge the patient-volunteers who agreed to wear our wireless device. Despite, I am certain, surplus burden in time of sickness. While your names will be forever kept silent, your efforts and contributions to science are known. No progress could have been made without you.

I am thankful for the many collaborators who made this research possible. Washington University in St. Louis – Medical School's scientists and physicians, Dr. Tom Bailey and Dr. Marin Kollef, from which I learned vital aspects of fusing technology and medicine. Our study coordinators: Emily Kuo and Pam Kemp, without whom I would have not survived a single day in the hospital (or find my way out of the complex hospital maze). Kevin Heard, from which I have received priceless support relating to the CDR integration. And our Artificial Intelligence partners Dr. Yixin Chen, and Yi Mao. There are at least a dozen more people whom I am neglecting to mention by name for brevity, but I am no less thankful for their support.

This research was made possible by Grant Number UL1 RR024992 from the National Center for Research Resources (NCRR), part of the National Institutes of Health (NIH) and NIH Roadmap for Medical Research. Additional fund ing was provided by the Barnes-Jewish Hospital Foundation and NSF through grants CNS-1035773 (CPS), CNS-1144552 (NeTS), and CNS-0708460 (CRI).

I am grateful to my academic mentor, Dr. Chenyang Lu, who insisted that I would embark on this project. It turned out he had more insight (or a time machine) than I did; as this turned out to be an important endeavor for my research career. Very special thanks goes to Greg Hackmann who has started this project and kindly graduated Dr. Gregory Hackmann (now with Google); and in doing so, allowed me to carry on and conclude.

Rahav Dor

*Washington University in St. Louis*

*August 2013*

Dedicated to my parents and grandparents.

This Masters Thesis provides me with very pertinent stage from which I can thank my parents, Dina and Lior Dor. Eema (mom in Hebrew) and Aba (father), you deserve a special recognition. Not because every Eema and Aba earns that for tolerating their adult offspring, from their teenage years until they decide that its time to become someone wordy of being called a man or a woman. Eema and Aba, in your particular case it is also, not just for being the best parents any child can wish for. You deserve my admiration and recognition for your unremitting support for the academic life in general, and my pursuit of happiness through the life of a researcher. I will love you forever and now its on paper.

For me this dedication will not be complete without saying what is always in my heart. Saba (grandpa) and Safta (grandma), on both sides, took agonizing punishment while saving themselves and their young children during the Holocaust. Menachem and Sara Doawer (individually) by leaving their families behind and having the courage and vision to go to Israel. Yosef and Lina Kasten by heroically escaping the ghetto and living in the forest for two and a half years. I admire the four of you. I wish that our family will never need to test our resolve, ever again. But if we do, we know what kind of spirit and blood flows through our veins. Good has triumphed evil. Our family and nation has won.

ABSTRACT OF THE THESIS

Difficulties and Opportunities in Building Resilient Clinical Monitoring Systems with Wireless Sensor Networks

by

Rahav Dor

Master of Science in Computer Engineering

Washington University in St. Louis, 2013

Research Advisor: Professor Chenyang Lu

Wireless Sensor Networks (WSN) can play an important role in improving patient care by collecting continuous vital signs and using it, in real-time, for clinical decision support. This Masters Thesis presents the architecture of, and our experiences with such a system. Our system encompasses portable wireless pulse oximeters, a wireless relay network spanning multiple hospital floors, and integration into the hospital Clinical Data Repository (CDR) database. We report our experience and lessons learned from a 14-month clinical trial of the system in six hospital wards of Barnes-Jewish Hospital in St. Louis, Missouri.

Such a deployment of a computer system around people daily lives makes it a Cyber Physical System (CPS) and exposes numerous challenges, where the advantages of WSN (for example being very low power) are in conflict with the unforgiving environment (for example overcoming interference in a heavily populated hospital, despite a low power transmission). As this Masters Thesis will show such a CPS can be made reliable, but there are numerous integration points in which the system can (and did) fail. Our experiences show the feasibility of achieving reliable vital sign collection, using a

wireless sensor network integrated with hospital IT infrastructure and procedures. We highlight technical and non-technical elements that pose challenges in a real-world hospital environment and provide guidelines for successful and efficient deployment of similar systems. At the end of this paper we hope that we will have given enough evidence and recommendations to convince the reader that the question "Can CPS that uses a WSN enjoy the advantages of very low power wireless communication, and be made reliable in a hospital environment?" is no longer open.

# Chapter 1

# Introduction

Early detection of clinical deterioration is essential to improving clinical outcome. Study has shown that 4 – 17% of patients will undergo cardiopulmonary or respiratory arrest while in the hospital as a result of clinical deterioration [1]. However, even following acute operations nurses measure the vital signs only 10 times during the first 24 hours [2], and much less vital sign taking is performed in less acute care hospital wards. Wireless sensor network (WSN) can enable real time detection of clinical deterioration by collecting continuous vital signs from patients, providing continuous coverage and alerts that could save lives. WSN also have the advantage that patients do not need to stay confined to their beds to benefit from monitoring. Compared to commercial telemetry systems, WSNs based on low-power network have advantages in both power efficiency and cost-effective deployment without fixed infrastructure [3]; and they are also an apt choice for field deployments.

While earlier deployment of WSNs in hospitals have shown promise [3, 4, 5, 6], they are not sufficient (including our own previous deployment) for clinical detection because their standalone WSNs are not integrated with the hospital's Clinical Data Repository (CDR[1]) and IT infrastructure and procedures – two key integration challenges that must be addressed in order to build clinical warning systems. Moreover, existing deployments tend to be limited to a single ward, while large-scale and long-term deployments spanning multiple clinical wards are needed to demonstrate the feasibility of WSN technology in hospitals. Such systems, deployed in close proximity people's habitats, for

[1] The CDR is typically used for data analysis, clinical research, and training and is separate from the production EMR (Electronic Medical Records), but it offers similar integration challenges in terms of hospital IT security procedures, and using the same data transport and messaging system for integrating between our system and the hospital system.

the benefit of people are Cyber Physical Systems (CPS). The continuous rub between people and the CPS surface challenges that must be addressed for the system to be reliable.

Our system and longitudinal clinical trial exercised these rub points and this Masters Thesis presents the architecture and deployment of our end-to-end clinical monitoring system. Our system integrates portable wireless pulse oximeters for patients, wireless relay networks spanning multiple hospital floors and wards, and departmental servers communicating with the hospital's Clinical Data Repository (CDR) database. This study is distinguished from the previous by the CDR integration and the study's scale. Exposing at least some of the rub points that similar CPS will encounter (we in fact suggest that any CPS deployment around people's daily lives will face similar challenges and could benefit from the recommendations we make). The system collects real-time vital sign data from the pulse-oximeters, transmits it to base stations over the wireless relay networks, and then feeds the data to CDR in real time, where it resides along with an historical corpus of the traditional clinical data. This enriched CDR data is then used to detect clinical deterioration using data mining techniques [1]. At the end of our clinical trial the system has been deployed for 14 months at Barnes-Jewish Hospital in St. Louis, with the first ward deployed in April 2011. The system spans three hospital floors, monitor patients in four floors, six wards, has between 14 to 16 wireless ”relay” nodes in each ward, four fixed ”base station” servers, and is integrated to the legacy hospital CDR systems.

Although our deployment's CDR integration, spatial scope, and duration may appear to have subtle differences from prior studies, these differences have had a surprisingly profound effect on the system's performance with many important design considerations for future deployments. The key lessons we learned are: (1) Multiple base stations can be used to increase the network reliability and resilience to hospital events. (2) Fault tolerant design is necessary to progress towards production-wordy CPS deployments. (3) We (re)discover the importance of end user indicators. (4) As part of the planning process, human factors must be addressed before the deployment starts.

(5) Hospital deployments must be treated fundamentally differently from lab (and even small scale, real) deployments, which typically focus on system issues. They must be engineered from the ground up to consider organization IT and hospital procedures.

The lessons learned during this deployment can serve as guidelines for future clinical deployments of WSN-based systems, and more generally to any large WSN deployment in the physical world where people presence is expected.

# Chapter 2

# Overview of the Clinical Trial

This fourteen month clinical trial is a joint research among Washington University School of Engineering, Washington University School of Medicine, and Barnes Jewish Hospital (BJH). The system deployed during this trial collects blood oxygenation and heart rate readings from consented patients every 1 minute. These readings are collected with the eventual goal of predicting patient deterioration by combining real-time vital signs of a patient with traditional clinical data of (similar) patients available in CDR.

The system discussed in this paper is part of a long-term effort to improve hospital care through prediction of imminent clinical deterioration and enabling physicians and nurses to render information-guided, and timely, attention to their patients. Our ultimate objective is to improve people's health, recovery during hospitalization, and reduce hospital readmission by extending the system reach to patient's homes. In order to achieve these goals we envision a novel two-tier system for clinical early warning [1]. The first tier uses data mining algorithms on existing hospital data records to identify patients at risk of clinical deterioration. Those identified patients will wear low-power, compact sensing devices that will collect and communicate real-time vital signs through a WSN. The second tier combines both the real-time sensor data and traditional clinical data to predict clinical deterioration, as well as suggest the most relevant clinical reasons based on data mining algorithms. This point is wordy of a short reminder for future studies. Even though the part of the system described in [1] showed that clinical events could be detected with accuracy of 0.9292 our physician collaborators wants more, and rightfully so. Because an early warning, rendered by the system before the onset of clinical conditions is detectable to the naked eye or by examination, is not practically useful. An expert system that will provide clues as for the possible reasons of the eminent deterioration is required. High-level description of the two-tier system architecture and preliminary results on the early warning, based on traditional clinical data (tier-1) was presented in [1], followed by an in-depth study of data mining algorithms for traditional clinical data presented in [7]. This Thesis is focused on the question weather it is realistic to build such systems with using WSN technology.

## 2.1 Patient Enrollment

In this sub-section we provide some details on the participant sample population and anecdotal field observations. A tier-1 data mining algorithm runs on the CDR database to identify patient at risk and alerts the study coordinator. The coordinator then tried to consent the patient to participate in the clinical trial. Table 2.1 presents the results of this enrollment process: Of the total of 355 patients tier-1 module triggered on, data was collected from 75 patients, another 12 consented but did not have data collected (this could occur, for example, when the device is not turned ON after placing it on a patient), and 42 refused consent. The remaining patients did not have the opportunity to consent due to their medical condition (being in contact isolation, not in their room when the study coordinator went to consent them, or were too sick to respond, and so on. See the table for more details). Beyond these 355 patients, an additional 173 patients were excluded from the table because they were discharged from the hospital before enrollment.

Table 2.1 Patient consent results.

| **Status** | **Count** | **Percent** |
|---|---|---|
| Consented | 75 | 21.1 |
| Consented, but not data was collected | 12 | 3.4 |
| Refused consent | 42 | 11.8 |
| Contact isolation | 131 | 36.9 |
| Not available or responsive | 67 | 18.9 |
| Do Not Resuscitate (DNR) | 15 | 4.2 |
| Not appropriate per nursing | 10 | 2.8 |
| Dementia | 3 | 0.9 |
| Total | 355 | 100.00 |

Tables 2.2 present the demographics of the 75 volunteers. The mean age of participants was 58 years.

Table 2.2 Demographics of the 75 consented patients.

| Race | Count |
|---|---|
| Black | 30 |
| Caucasian | 40 |
| Other | 5 |
| Total | 75 |

| Gender | Count |
|---|---|
| Female | 51 |
| Male | 24 |
| Total | 75 |

## 2.2 Field Feedback

At the end of the trial we asked the study coordinator, who was responsible for consenting patients and applying patient devices, for her observations. We were interested in the participants' motivations, complaints, and experience wearing the device. These anecdotes are paraphrased below.

> Patients volunteered due to their awareness of benefits to future treatment. Most complaints were regarding the inconvenience of the device, possibly due to other equipment they wore. Importantly it centered around mobility and their ability to perform ADL's (activity of daily living). Participation of some subjects was short, in part due to a short hospital stay, but we had a few patients with long term participation (three days) without any problems or complications. The majority of issues required going back to the participant's room to investigate the patient device.

There are many factors that affect how long a volunteer would give up her or his time and comfort to advance science, in particular in time of illness. Figure 4.5 plots the time each patient wore our device and a histogram of the data. The left figure shows that participation time did not improve as our clinical trial progressed. This is contrary to what we expected because over the life of the trial we have introduced posters to advertise the study and its importance. We have also added a short promotion pep talk given by the study coordinator when she consented volunteers.

The histogram on the right of Figure 2.1 shows that the two greatest groups of volunteers either wore the device less than 10 hours or between 20 to 30 hours.

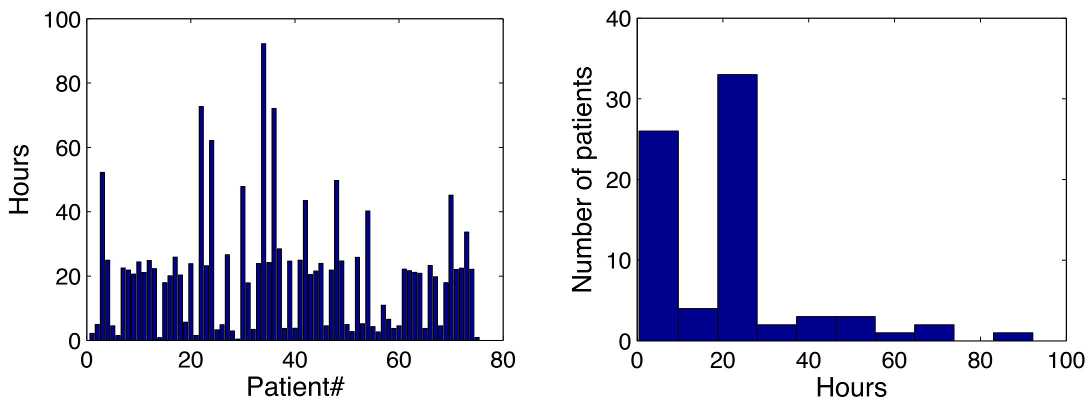


Figure 2.1 A plot of the time each patient wore the device through the trial duration (ordered by date), and a histogram showing the distribution of this time.

# Chapter 3

# System Architecture

In this section we discuss the architecture of the clinical monitoring system and debate which type of wireless network would be appropriate in a clinical setting. The system consists of three major components. Patients wear a wireless patient device that collects vital sign data (specifically, blood oxygenation and pulse data) during their hospital stay. As vitals are collected, the patient device transmits the data to a wireless relay network deployed throughout the hospital wards where our clinical trial is running. The relays form a low-power wireless mesh network that carries the vital sign data from the patient devices to our servers (referred to as base stations). The base stations act as an interface between the WSN and the hospital's IT systems (via wired Ethernet) and feed real-time vital sign data to the CDR database. The base stations also provide network monitoring and logging services that have been invaluable to system performance evaluation during our clinical trial.

The patient device and wireless relay networks reuse the same design from an earlier clinical trial detailed in [3], while the base station software has been extended substantially with new services to support the CDR integration and network monitoring. The key distinctions between this study and the earlier study reported in [3] are three-fold: (1) The integration with the hospital IT infrastructure and CDR compared to standalone WSN. (2) A large-scale, multi-floor and multi-ward deployment compared to single-ward deployment. (3) Ours is a longitudinal study (14 months). Each of these characteristics led to important new architecture and code-based for our system, as well as operational challenges and insights – as reported in this paper.

## 3.1 Patient Device

Each patient device integrates a TelosB [8] mote running the TinyOS [9] operating system, with a Nellcor pulse oximeter sensor [10], attached to the patient's finger to collect pulse rate and blood oxygenation level (SpO2). Nellcor's sensors use color-coded connectors; the appliance developed for the first study is compatible with the purple coded sensors. Nellcor offers a number of sensors, prescribed according to the patient body weight, duration of monitoring, and activity level. Some sensor models are disposable and some reusable. We were interested in long participation periods and as a result we were attentive to the connivance of having sensors being attached to the patient for long durations and their reliability due to movement. Therefore, as part of this research we performed a mini-study with and determined that the finger sensor model MAX-A (adhesive, disposable) was the most convenient both for placement as well as wearing. This is the only model that we ended up using throughout the trial. The -A denotes a short cable, which suited our patients, but a longer -AL model is also available. Other models that could be use are the SC-A that uses Velcro instead of an adhesive and should be used when the patient has fragile skin. This model also utilizes a high-efficiency LED to enhance pulse acquisition for patient with thicker, darkly pigmented skin, or weak pulses. We have never gotten around to actually using this model because patient's characteristics were not available when the system alerted the study coordinator with a candidate patient. In retrospect this was unfortunate because we have lost some patient data exactly to those reasons. The Dura-Y (an ear sensor that can also be placed on the finger) sensor, which was recommended by BJC nurses, was not used in our trial because it was harder to apply and frequently fell off. For research and development in the Engineering lab we used the reusable DS100A-1 model because it was very convenient to quickly apply to oneself. Figure 3.1 provides a photo of the patient device and a data flow diagram of the vital signs in the patient device.

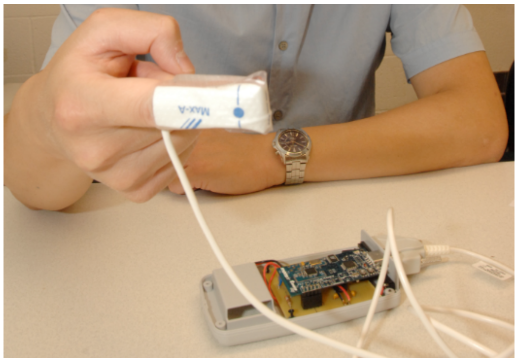

Figure 3.1 An open patient device with a MAX-A sensor attached is shown on the left.
The data flow diagram below shows, from left to right, the vital signs path through the device. The Nellcor sensor acquires heart rate and blood oxygenation and the OxiLink interface interprets the light absorption to digital data. The interface board connects the OxiLink RS232 output with a TelosB Mote.

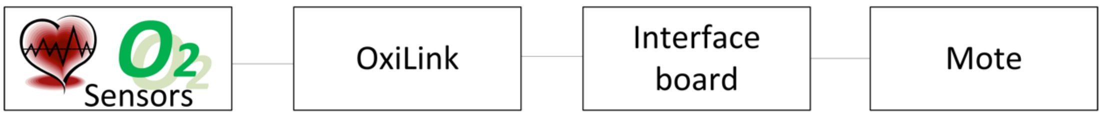


The sensor operates by emitting red and infrared light through the body to a photodetector. The photodetector's output is wired to an interface called OxiLink made by Smith Medical. interface [11], which interprets the light absorption data to provide meaningful heart rate and blood oxygenation readings. A custom circuit board built for the previous study is used to interface the OxiLink RS232 output with the TelosB and to perform voltage conversions needed to physically interface the device's constituent parts. A 9V battery powers the entire patient appliance. The software we authored for the TelosB controls the power to the sensor and thus collects data at a rate of 1 reading per minute; otherwise it shuts down the power to conserve battery. After receiving a sensor reading the TelosB routes the data to the relay infrastructure discussed in the following subsection using its onboard CC2420 [12] radio. The patient's device locate relays using the Dynamic Relay Association Protocol (DRAP) [13], a one-hop protocol specifically designed to locate a larger routing infrastructure under mobility.

## 3.2 Relays Network

WSN network coverage is achieved by plugging TelosB nodes into electrical outlets in patients' rooms, offices, and elsewhere where it seemed spatially reasonable. The relay devices use their onboard radios to form a mesh network based on the IEEE 802.15.4

low-power wireless networking standard [14]. Figure 3.2 presents a high-level carton of the resulting network topology. The Collection Tree Protocol (CTP) [15] is employed on the relay nodes to establish data collection trees rooted at a base station. 802.15.4 wireless links are highly dynamic, due to factors such as external interference. Hence, CTP is a dynamic routing protocol that uses periodic beacon messages to maintain the routing infrastructure, autonomously changing the links used to route packets as their reliability varies over time.

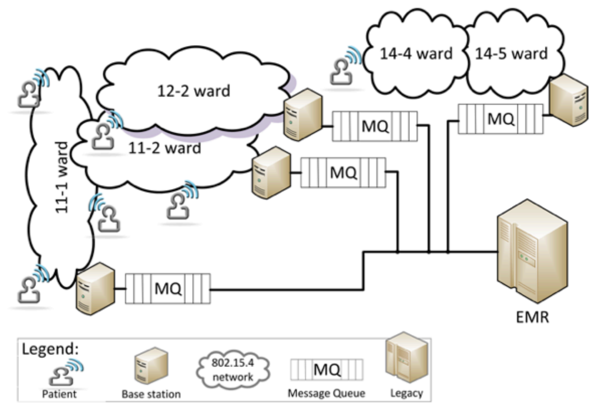


Figure 3.2 A high-level view of the network topology. The figure shows the ward floor and number, e.g. 11-1 is ward 1 on the 11th floor. Each WSN cloud is attached to a single base station at a certain time, but the sensors participating in each cloud do dynamically change cloud membership in response to changing environmental conditions, and each cloud can be rooted in a different base station for the same reasons. Each base station interface to the legacy hospital CDR via MQ Series (by IBM) that transports data between both ends.

In addition to patient data, relay nodes send “beacon” packets to a base station at a rate of one packet per minute. These beacons provide diagnostic information about the network’s operation even when no patients are enrolled.

A subtle but important feature of CTP is that each relay aims to find a high-quality path to some base station, not any base station in particular. Multiple base stations may be deployed in the same network to provide redundancy, a feature we exploited in our deployment. As we will discuss later in this paper this decision had a profound (and sometimes surprising) impact on our network’s behavior.

Figure 3.3 presents a snapshot of our relay deployment, taken at a point when our network was deployed in five wards on three hospital floors (11, 12, and 14). Each ward was equipped with 14–16 relay nodes and one base station. Black arrows indicate the links used to form the network's routing trees. As discussed above the routing paths are highly dynamic and thus this snapshot frequently changed. Light purple circles denote the infrastructure relays. Darker purple denote base stations. Red shows relays that are currently disconnected from the base stations. The bottom left side of the image shows the network on 11th floor and its two base stations, number 50 on the bottom left and 52 on the center right. The bottom right part of the figure shows the network on the 12th floor and its base station 60. The top of the image shows the 14th floor and its base station 73.

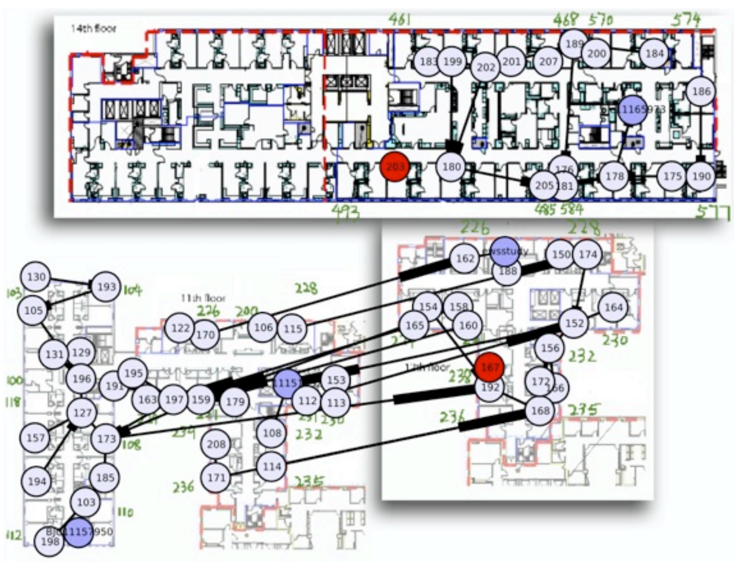


Figure 3.3 Relays infrastructure across five hospital wards.

We expected that the respective wards would generally form their own isolated routing paths and that connections among wards would be rare. Nevertheless we placed all the relays on the same wireless channel (allowing any relay to theoretically connect to any base station) and deployed several relays in the corridors between the two wards on the 11th floor. Our intention was to provide resilience against base station failure by allowing relays in one ward to send their data to the base station in another ward, in case of a need. However, as shown in figure 3.3, the relays routinely formed links across wards. Moreover, because the 12th floor deployment was directly above the 11th floor, routing trees often crossed through the ceiling and floor despite (as we've been told) each floor being a tick slab of concrete, metal, and other infrastructure). We also discovered that the 11th floor infrastructure could reliably collect data from patients in the corresponding 10th floor wards, even without any relays installed on the 10th floor. We discuss these observations later in the paper. The 14th floor deployment was located two floors above and in a different section of the building and remained isolated for the duration of the deployment.

## 3.3 Base Stations

Each base station consists of a laptop with USB-attached TelosB mote, and acts as an interface between the 802.15.4 network and the hospital's wired Ethernet infrastructure. The base stations are Intel based machines running the Microsoft Windows XP operating system. The base station manager software is authored in Python and Java (two languages where used for reasons explained below). We installed a SSH Server on the machines to allow us to remotely administer them. Each base station is deployed with base station manager software that we authored with the core functionality depicted in Figure 3.4. As vital signs produced by patients arrive from the relays mesh network to the mote attached to one of the base stations, the base station software consumes the vital sign data.

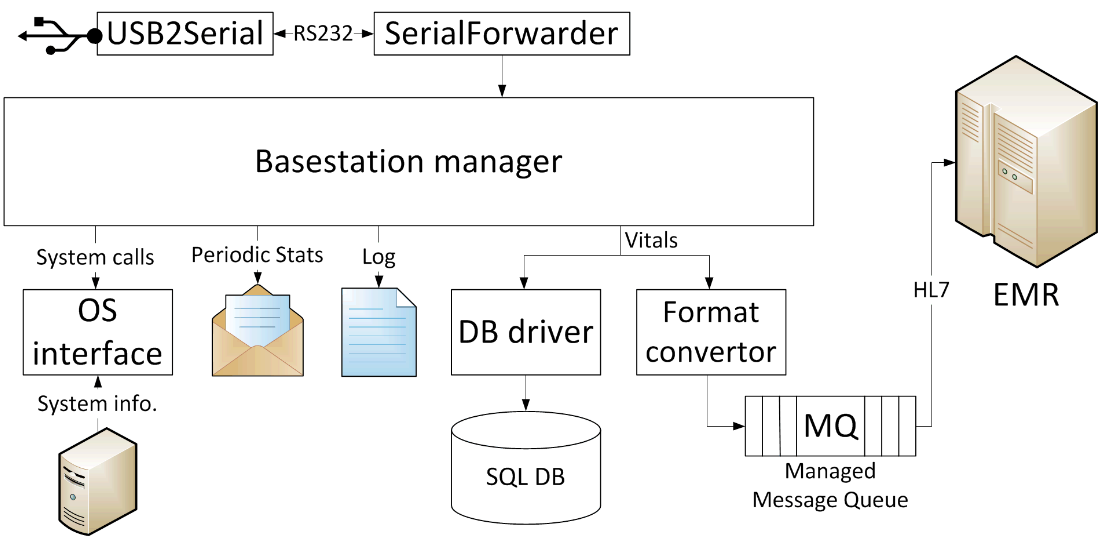


Figure 3.4 Data flow through the data collection system.

Data arrives from the wireless network to the mote attached to the USB port of a base station. A module called *Raw2Obj* (not in the figure), and *SerialForwarder* (which is part of TinyOS tool-chain) convert the raw data into packets. *Raw2Obj* was not developed by this project. It is part of a software library called *WSNtestBed* developed at our lab for running our motes test bed. We note that *SerialForwarder* could assume *Raw2Object* functionality now that Jython is an integral part of the software stack. We added *SerialForwarder* to our architecture because instead of directly connecting (and blocking) to the USB port, *SerialForwarder* allows multiple applications to consume the data by providing a TCP socket. Our software, the *Basestation* manager, is one of the consumers of this socket.

Two types of packets arrive to each base station: network status (beacons") and patients' vital signs. Both are saved in a local *SQL DB*; with the vital signs also formatted as HL7 messages, which are then sent to the hospital *CDR* via a message queue *MQ*. The other data paths of the system (System calls, Periodic stats, and Log) are described in the next sub-sections.

### 3.3.1 Software Stack

The base station software is an evolution of software (first version was authored in Python) we used in prior studies. Predominately, the addition of Java was driven by the need to use a Java based MQ library because the Python based MQ was too costly for our research budget. Data is exchanged between Python and Java classes was made possible by using Jython [16], which is an implementation of the Python programming language in Java. As such, it provides a seamless objects exchange between the two languages. This functionality is denoted in Figure 3.5 as Native Python/Java interaction layer. Furthermore, we purposefully chose an ODBC driver provided by Jython. Thus, through this Python/Java interaction layer, Python modules running on the base station can talk to the DB directly, and since the DB driver is authored in Java it makes the DB natively accessible for the Java modules as well.

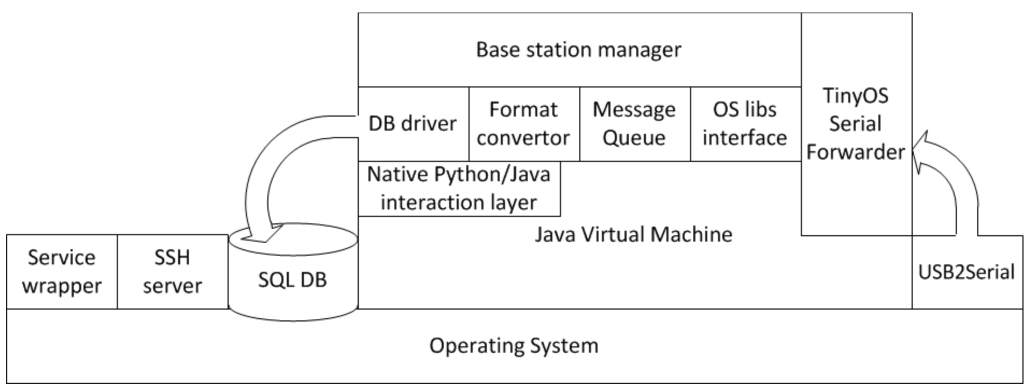


Figure 3.5 Base station software stack (major components)

*Service wrapper* allows our system to run as an OS service (always working and available whenever the OS is available). The *SSH server* allows us remote access for system administration; Depending on the circumstances, we sometimes found that either *SSH* or *Windows Remote Desktop* were more convenient. Surprisingly there were cases where only one of them gave us access, which of course occurred when we needed it most, so having both saved us a few trips to the hospital. The *OS libs interfaces* allowed us to query the OS, DB, and machine.

### 3.3.2 Health Level 7

The vital signs arriving to the *Base station manager* are formatted according to the international Health Level Seven (HL7) standard [19] with a converter built on top the HAPI library [20]. The module named *HL7 Formatter* in the Figure 3.4 depicts this functionality. The HL7-formatted vital sign data are then inserted to a Message Queue (*MQ*) that reliably transports it to the hospital CDR, over the hospital's wired Ethernet network. This particular deployment used the IBM WebSphere MQ library [21] for consistency with existing hospital IT systems.

### 3.3.3 SQL Database

The vital signs the network activity are written to the local SQL DB. The base station manager also logs any important event occurring on the system itself to a local sequential file. In this implementation we have used the PostgreSQL database [18] version 9.x.

The DB consist of 2 active tables:

- paths – a time stamped log of the path via which each node was connected to a base station. Every 1 minute (assuming that a node had a path to a base station at that point in time) each node path is recorder.
- vitals – in which the patients vital signs are stored

The *vitals* tables stores **all** data messages received from patients, including invalid readings. To filter erroneous messages queries on the *vitals* table should include (among other filters needed for your desired query) the clause:

```
where systemFlags = 0 and signalLevel >= 2 and avgSpO2 != 127 and
heartRate != 255
```

The data base also has 4 tables (receive_log, send_log, stats, timing) that are leftovers from the prior study and were kept in our project only to for compatibility with the existing code base.

When a consolidated view of the data was desired the databases from the four base stations where replicated to my laptop and a merged database for analysis was created. Databases DB50-End, DB52-End, DB60-End, where DB## designate the base station ID and End designate data at the End of the trial are available. The consolidated database is named RDS_Analysis-End. This database also contains an additional table named vitalsFromKevin that holds vitals signs data as they were stored in the legacy CDR.

The *paths* table in the consolidated database was cleaned from erroneous records using the following query, and the analysis presented in this paper was done against the *patch_c* table:

```
create table paths_c      -- c = clean
as select * from paths
   where (
       (node between 100 and 180)   -- skip 181, has only 1 record
    or (node between 182 and 187)   -- 188, sends multiple packets
                                       per minute, probably due to
                                       hardware failure
    or (node between 189 and 198)   -- 199, has only 1 record
    or (node between 200 and 216)   -- currently last node ID in the
                                       deployment is 208, our
                                       temporary mote experiments
                                       starts at 217
    )
```

When using the vitals signs data as it appears in the hospital CDR, the *vitalsFromKevin* tables should be used. This table designates missing vital signs as records with status -1 (a patient is suppose to have a vital sign reading every minute). In queries where only the received vital signs data is of interest use the following query (among other filters needed for your desired query) to filter those out:

```
where not ("O2_status" = -1 or "pulse_status" = -1)
```

### 3.3.4 Continuous Operation

Because vital signs arrive in an asynchronous and continuous manner a malfunctioning base station means loss of data. To allow the base station software to operate continuously (for example following loss of power or mandatory operating system restarts after installing patches) we used a Service Wrapper that enables Java applications to be run as a Windows Service or UNIX Daemon. In this implementation we used the Tanuki [17] Java Service Wrapper. Tanuki also provides tools to monitors the health of the application and the JVM it run on. Running as an OS service starts our application as part of every Windows boot process.

### 3.3.5 Network Monitoring

An important feature of our base station software is a suite of network monitoring services that have proven invaluable for the maintenance. For diagnostic and debugging purposes, each base station is equipped with a database server (in this implementation PostgreSQL) that is used to locally log all messages received from the relay network (in the *paths* table). This log includes the relays' diagnostic beacons, which are not forwarded to the hospital CDR. Extraordinary events are also logged in a time-stamped sequential file for easy access.

### 3.3.6 System's Vital Signs

To track the base stations' health each server produces and emails a nightly report every 24 hours (configurable). In addition to the information contained in each report the email itself has proven to be a useful diagnostic tool during the trial: failing to receive a report is a good early indicator that a base station is disconnected or otherwise malfunctioning. The key sections of this automatically generated and emailed report are described in the Figures 3.6 and 3.7. Among other items each report lists the patient devices that sent data to this base station in the past 24 hours, the number of vital sign readings for each device, a list of relays and their current path to the base station, the status of the local PostgreSQL database, network connections, the status of the attached mote, and the state of other hardware sub-systems such as storage.

**Network Status**

| Interface Name | MAC | IP |
|---|---|---|
| 11a/b/g Wireless LAN Mini PCI Express Adapter - Packet Scheduler Miniport | 00:23:4D:21:63:2D | |
| Intel(R) 82566MM Gigabit Network Connection - Packet Scheduler Miniport | 00:1F:E2:18:86:A0 | /10.33.146.144 |

**Mote Status**

| Name | ID | Location |
|---|---|---|
| COM5 | XBTFV8VJ | Crossbow Telos Rev.B |

**Infrastructure Status**

| Mote ID | Last beacon | Path |
|---|---|---|
| 103 | 2012-02-29 23:59:41.644000 | 103 |
| 127 | 2012-02-29 23:58:21.926000 | 127, 157, 194, 103 |
| 129 | 2012-02-29 23:59:30.783999 | 129, 127, 157, 194, 103 |
| 130 | 2012-02-29 23:55:21.460999 | 130, 196 |
| 157 | 2012-02-29 23:59:11.487999 | 157, 194, 103 |
| 173 | 2012-02-29 23:59:52.878000 | 173, 194, 103 |
| 193 | 2012-02-29 23:55:31.585999 | 193, 130, 196 |
| 194 | 2012-02-29 23:59:42.519000 | 194, 103 |
| 196 | 2012-02-29 23:55:18.132999 | 196 |

Figure 3.6 The Network Interface (access to the hospital Ethernet) and attached Mote sections of the periodic status report; as well as all of the relays that are in this base station's cloud and their path to the base station.

**Power Status**

| | |
|---|---|
| ACLineStatusText | Online |
| BatteryFlagText | High—the battery capacity is at more than 66 percent |
| BatteryLifePercent | 88 |
| BatteryLifeTime | unknown |
| BatteryFullLifeTime | unknown |

**Database Status**

| | |
|---|---|
| PID | 4248 |
| Database size | 320M |
| Open cursors | 1 |

**Available Resources**

| | |
|---|---|
| Free memory | 16941K |
| Total memory | 54312K |
| Max memory | 253440K |
| Usable disk space | 43752M |
| Free disk space | 43752M |
| Total disk space | 66297M |

Figure 3.7 Since our base stations were based on laptops, which were often disconnected from a power source, it became important to know if it is connected to a power source (ACLineStatusText) and the expected time the battery has left. Also reported are the status of the DB and other general resources of the machine.

Most of these pieces of information require access to the local OS service libraries. We have achieved this functionality using a Java library called Java Native Access (JNA) [22].

To help us understand the network performance we developed a utility that queries all base stations *paths* tables, each base station at a time, and draws map the network topology of the last 15 minutes. Figure 3.3 presents an example topology map produced by this utility. These maps have been particularly useful for identifying network failures caused by equipment disconnections, which we discuss in more detail in section 4.1.

## 3.4 Network Type

Any system's architecture, working in a certain cyber and physical environments, would have advantages and drawbacks owing to the type of wireless network in use. As any design alternatives we believe that such questions should not be religiously considered, that the preferred network will need to be evaluated between major releases, and that even for a given point in time there is more than a single good choice. In this section we offer our design considerations as they were in the beginning of the research endeavor, along with some alternatives that become more relevant as the trial ended.

Unsurprisingly, WLAN (IEEE 802.11, or informally WiFi) [30] seems to offer an advantage because it was already installed in the hospital. However in closer analysis we found out that while that WiFi network is available almost everywhere in the hospital, reception is sporadic in different sections of the hospital. We determined that in order to provide good network coverage across a hospital ward, in particular due to the hospital environments, there is going to need to be an access point every 20 to 30 meters. Comparing to a WSN relay that we planned to install every 10 to 20 meters (before we added extra relays as discussed later in this document), then, a WiFi will save between 40 to 67 percent of relays. However because each WiFi access point requires an Ethernet jack, the total cost of ownership (TCO) will be higher if WiFi is used. It is worth noting that we deliberately chose to use only the printed circuit board antennas of the relays. Had we used external antennas then the number of WiFi access points and WSN relays would be equivalent.

When considering WiFi one also need to take into account that the network may experience irregular congestion periods because it is not known how many guests will be using it during the day. There are of course medical grade (and frequency) WiFi networks, but BJC did not have these available; installing them would require the same amount of infrastructure as a regular WiFi and we wanted to offer a clinical system that could be deployed with minimal infrastructure.

WSN offers an advantage in term of energy consumption. Typically WSN transmission power ratio is between -24 to 0 dBm, while WiFi is 20 to 30 dBm. This is an power saving factor of 1,000 to 25,000 in favor of WSN.

An additional WSN advantage is the wakeup time, which is 1 to 10 ms for WSN and about 30 ms for WiFi. In a network of patient device that are mostly asleep and wake up, at most, every minute – it is an advantage in both response time and energy used during the wakeup period.

We did not feel that a network technology such as any one of the standards (GSM, CDMA) available for cell phone networks due to the short distance between the patient devices to the base stations. However a completely different architecture, which connects a sensor to a smart phone provided to the patient or allocated to the room should be considered. The sensor(s) could communicate with the phone using Blue Tooth or similar network, and the phone will then serves as a mini, or full pledge, base station, connecting to the rest of the system. We could not explore this approach because it was not part of the original grant and such architecture has budget implications.

For completion we would like to mention three additional technologies that should be considered in future deployments. These are Roofnet network by MIT [31], 2net by Qualcomm [32], and white space communication [33].

We choose for our system the WSN (802.15.4) network because it could be deployed with little to no infrastructure, one simply need to plug a mote into a power outlet. A patient device wakes up every minute for a few seconds and then goes to sleep, therefore wakeup time was important to us. And we wanted our device to work at least a week with a single 9V battery. With those choices, in the beginning of the system development, WSN was almost the only option available.

# Chapter 4

# System Reliability

In this section, we analyze our deployment's network and system reliability. We first assess the overall network reliability based on the beacon packets sent by each relay. Recall that each relay periodically sends a beacon packet to the base stations using the same CTP protocol used for transporting regular vital sign data. The collected beacons give us a comprehensive view of reliability in every area of the network even when there is no patient enrolled. We then provide an analysis of the reliability of vital-signs-delivery from enrolled patients, offering a direct view of our CPS ability to deliver real-time vital signs using a WSN. This reliability analysis shows that despite many improvements (articulated in chapter 5) that can still be done to our system, this WSN was already reliable enough despite our faults.

Due to its location, the 14th floor formed its own network separate from the rest of the network we study here. As this single-ward network is similar to our previous deployment studied in [3], henceforth our analysis focuses on the challenges posed by the larger, multi-ward deployment on the $10^{th}$ to the $12^{th}$ floors.

## 4.1 Relay Connectivity

To analyze the relays availability over time we define a relay as connected based on the fraction of its beacon messages that were successfully delivered to at least one of the base stations. On each day, for each relay, we calculate the ratio of the number of delivered beacons to the expected 1440 beacons/day (1 beacon/minute). We note that a beacon message may be received by any one of the base stations due to CTP protocol's design. Our DB is distributed over a number of base stations and we will need to combine the queries from all of them for correct view of reliability. We define the Packet Delivery Ratio (PDR) threshold to designate relays as connected or disconnected each day as follows:

> To be connected, a relay must have successfully delivered 90% of its beacons to the base stations that day.

Figure 4.1 plots the number of relays connected to some base station over the course of the deployment. The red curve indicates the number of connected relays for each day of the deployment, while the purple horizontal lines indicate the total number of relays physically deployed in the network; shaded regions mark significant events during the deployment that we discuss below. We note that Figure 4.1 begins at May 24, 2011: due to factors described in Section 5.5 (Impact of IT procedures) the diagnostic logs needed to produce this graph are partially incomplete between the initial deployment and May 23; an unfortunate circumstance that mandated that we will exclude this period from our analysis.

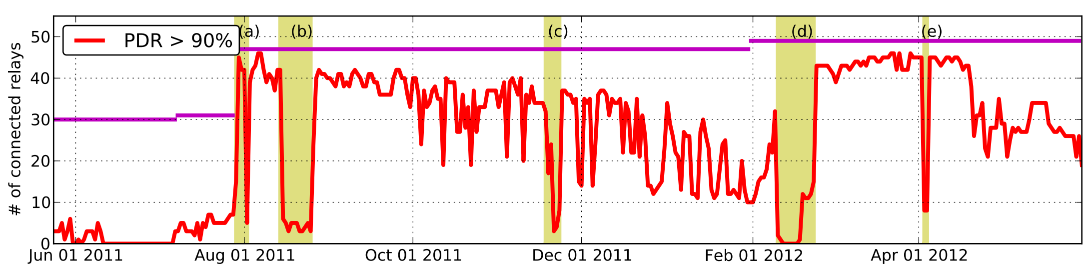


Figure 4.1 The network reliability from May 2011 to May 2012.

The purple line marks the number of deployed relays. The red plot represents the number of relays connected to the base stations. Significant events are shaded: (a) The deployment was expanded to ward 3 on July 29. (b) Between August 14 to 25 base station 52 was unplugged from power. (c) Between Nov. 18 to 23 base station 52 did not work. (d) Between Feb. 10 to 23 relays disappeared again and for a while we were prevented from replenishing them because they were in patient isolation rooms. This caused the three available base stations to receive only a fraction of their regular traffic. (e) On April 3 and 4 base station 52 received about half of its usual traffic for unknown reason; the problem resolved itself without intervention. At the end of the trial relays were slowly disconnected without replenishing, leading to the gradual decrease in reliability.

Starting in April 2011, the network was deployed with 11 relays and two base stations (denoted 50 and 52) at two corresponding ends of wards 1 and 2. When we saw the poor network reliability we added 19 additional relays a few days later. Even though we now had good relay density, the network reliability remained poor. The reliability stayed unacceptable through July and we initially attributed this to relays that physically disappeared from the hospital, and base station issues which are discussed in the next section. But despite relay replenishment, our work on improving the base stations, and other improvements we had made to the system, it did not reach our expected level of reliability. On July 29, designated as event (a) in figure 4.1, the deployment was expanded, adding ward 3 exactly one floor above ward 2. Ward 3 was furnished with its own base station (denoted 60) and 16 additional relays; since ward 3 was within radio range of wards 1 and 2, the new relays and base stations immediately merged into the existing infrastructure.

Admittedly more research into the affect of people on low energy networks is needed, but we hypothesize that the presence of people in this busy hospital, which treats about 400 new patients each month, had a profound effect on the network performance. The deployment of ward 3 provided paths to base stations through the ceiling and floor, routes that to some extent avoided links through corridors that are may be heavily populated with people.

During the initial deployment, we deliberately allowed ward 1 and 2 networks to merge, expecting that the redundancy would improve connectivity. However, even before significant relay disappearances took place (discussed in the next section) the network struggled to achieve high reliability before August. We also noticed anecdotally that the network tended to route data to ward 2's base station; even relays in ward 1 would often route data to ward 2 as opposed to the (geographically closer) base station in ward 1. This observation is also reflected in figure 4.1: when ward 1 experienced an outage from July 22 to July 25 (not highlighted as an event in the figure), it had little effect on network connectivity: the number of connected (PDR ≥ 90%) relays remained low, but was largely unaffected by the outage. After the expansion ward 2's base station has also proved to be quite important: reliability fell greatly during event (b) between August 14 and 25, when this base station was disconnected from power. Looking closely at the data from during this outage, we found that it even affected the recently deployed relays in ward 3. From these observations, we concluded that the two floors' networks strongly depended on each other for high reliability.

Even though our deployment planned for some cross-ward traffic this phenomenon was much more prominent over the duration of the trial than we anticipated. Figure 4.2 shows the proportion of time each relay was connected to each base station, with the least favored base station (50) handling 13.4% of the traffic. It is worth noting that base station 50 is installed in the physicians' consultation room, which has continuous foot traffic — possibly another indication to the interaction between people and WSNs. Remedies to these conditions is offered in the next section.

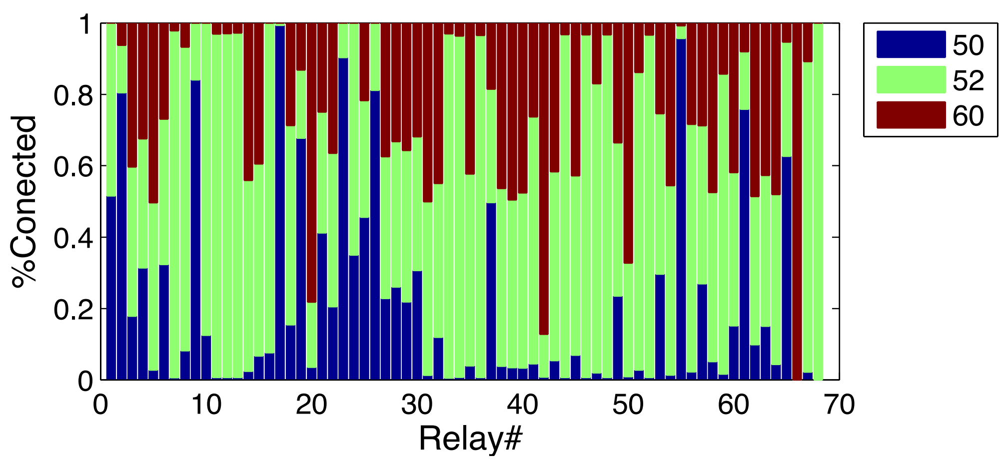


Figure 4.2 The percent of time relays were connected to a base station (50, 52, and 60). We see that almost every relay in the deployment divided its time among multiple base stations.

# 4.2 3D Networks

Many of the motes used routes to base stations that are not near them. See for example the odd choice that relay #115 is making in Figure 3.3. Instead of connecting to the nearby base station 52 (the purple circle that just below relay 115 in the figure or is about 7 meters away across the corridor) it connects through the ceiling to relay 174, then to 152, which in turn connects to base station 52 through the floor. When we query the DB, it turns out that relay #115 required paths via the ceiling (and possible back through the floor) 45% of the time. Many other motes, in fact most of them, required such paths (see Figure 4.3). We name such network topology, 3D Networks.

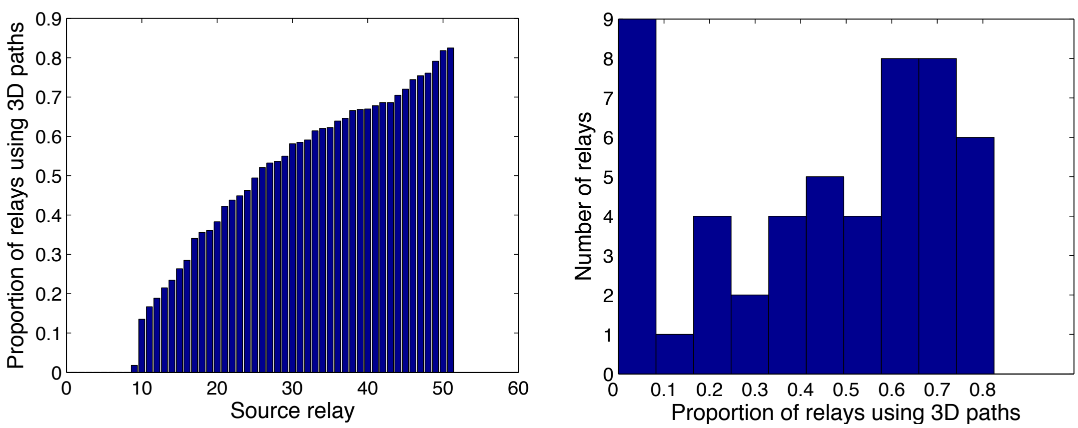


Figure 4.3 The proportion-of-time motes on the 11$^{th}$ floor used routing paths via the ceiling (floor). The left plot is ordered by proportion and we see that 9 motes never required such paths.
The right plot is the histogram of the data, demonstrating that most motes require 3D paths in high proportion of their routes.

Figure 4.3 show that 9 relays (number 102, 110, 111, 118, 119, 125, 126, 128, and 177) did not use 3D paths. Relay 177 is the only exception in this group, using 3D path 1.8% of the time. Those 9 relays disappeared from the hospital before the deployment of the network on the 12th floor. Had they been part of the network following the 12$^{th}$ floor deployment they might have used this 3D path diversity as well.

Possibly more interesting investigation would reason on relays requiring 3D paths in more than 50% of the instances. To explain our analysis we draw your attention to Figure 3.3, and Tables 4.1 displaying all relays requiring 3D paths in more than 50% of the instances, and 4.2 requiring 3D paths in less than 50% of the instances. The tables list the relay#, the proportion of the instances it used 3D paths, whether a direct path to a nearby base station required an airway path through areas that may be with heavily populated, and a comments column.

Table 4.1 Relays using 3D paths more than 50% of the instances. The table is sorted by the 3D path% column.

| **Relay#** | **3D path%** | **Via people** | **Comments** |
|---|---|---|---|
| 163 | 82% | Yes | |
| 151 | 82% | Yes | Family waiting room |
| 208 | 79% | Yes | |
| 191 | 76% | Yes | Family waiting room |
| 197 | 75% | Yes | |
| 104 | 74% | Yes | Family waiting room |
| 195 | 72% | Yes | Family waiting room |
| 124 | 70% | Yes | Family waiting room |
| 169 | 69% | Yes | Elevator lobby |
| 122 | 69% | Yes | Linen closet |
| 171 | 68% | Yes | |
| 196 | 67% | Yes | Elevator lobby |
| 193 | 67% | Yes | |
| 131 | 67% | Yes | Nurses' station |
| 129 | 65% | Yes | Lead Charge Nurse office |
| 159 | 64% | Yes | |
| 117 | 62% | Yes | |
| 130 | 62% | Yes | Nurses' station |
| 170 | 61% | Yes | |
| 185 | 59% | Yes | |
| 132 | 59% | Yes | |
| 173 | 58% | Yes | |
| 107 | 55% | Yes | |
| 120 | 54% | Yes | |
| 105 | 53% | Yes | |
| 161 | 52% | Yes | |

Table 4.1 demonstrate that all relays that used 3D paths in more than 50% of the instances, did that despite sometimes having a more direct path to a nearby base station. The "Via people" column suggests that all of these plausibly more direct airways had to pass through areas in the hospital that, at times, are heavily populated. To the extent that at a given time this people/radiation interaction rendered a direct path less reliable the relay chose an alternate 3D path. In the "Comments" column we also offer additional support for this hypothesis. We see that there is a direct correspondence between the proportions of instances that a 3D path was required to the areas that are more likely to have larger amount of people.

The next Table, 4.2, demonstrates that relays, which used 3D paths in less than 50% of the instances, were installed in areas with likely less dense population. The three exceptions of relays in this table that did require a passage through people are of relays that are 1 hop away from a nearby base station. The nearness to the base station is the probable cause compensating for the energy absorb by people and the reason that these relays appear in table 4.2 instead of 4.1.

Table 4.2 Relays using 3D paths less than 50% of the instances. The table is sorted by the 3D path% column and the rows of relays with a routing path to a base station that require passage through heavily populated areas are shaded.

| Relay# | 3D path% | Via people | Comments |
|---|---|---|---|
| 114 | 49% | No | |
| 123 | 46% | No | |
| 115 | 45% | Yes | 1 hop away from BS #52 |
| 127 | 44% | No | Staff pantry |
| 157 | 42% | No | |
| 179 | 38% | Yes | 1 hop away from BS #52 |
| 108 | 36% | No | |
| 198 | 36% | No | Physician workroom |
| 106 | 34% | Yes | 1 hop away from BS #52 |
| 113 | 28% | No | |
| 121 | 26% | No | |
| 153 | 23% | No | |
| 194 | 21% | No | |
| 112 | 19% | No | 1 hop away from BS #52 |
| 103 | 17% | No | 1 hop away from BS #50 |
| 109 | 14% | ? | Location uncertain |
| 177 | 2% | No | |

An interesting question is weather there was something special about the 11$^{th}$ floor that caused this behavior, or is the 12$^{th}$ floor relays going to behave in a similar way? If the 11$^{th}$ floor motes seek the 12$^{th}$ floor infrastructure in such high proportion, will the 12$^{th}$ floor relays do without 3D paths?

As an answer we provide Figure 4.3 showing the behavior of the 12$^{th}$ floor relays. Indeed, we see that most of the relays (12 out of 15) required 3D paths in more than 50% of the routes.

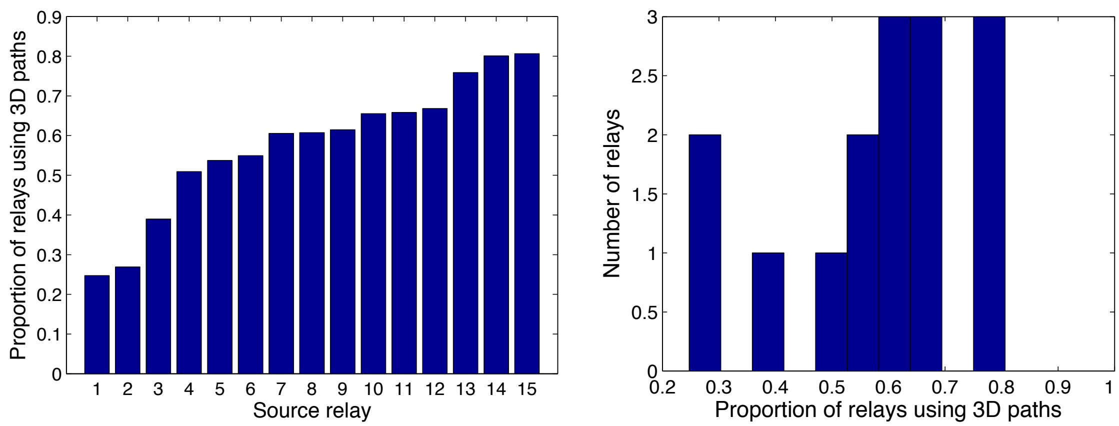


Figure 4.4 The proportion-of-time motes on the 12th floor used routing paths via the floor (ceiling). The left plot is ordered by proportion.
The right plot is the histogram of the data, demonstrating that 12 relays out of 15 required 3D paths in more than 50% of their routes.

Since we did not solidify our 3D Networks idea during the trial and the analysis we perform here is retrospective, further research will be needed to validate this hypothesis, and there may be other causes, beyond people for this behavior. It was also suggested by our Medical school collaborators that hospital food carts might pose a periodic daily problem. These metal carts are high and could act as Faraday cages, but they are not present in the corridors long enough (time wise) to explain such high proportion of 3D paths.

Before the deployment in the 12th floor this need for 3D Networks critically affected our system performance. It is likely to affect similar CPS deployed around people's daily lives. WSNs clearly offer advantages with its low energy consumption, duty cycling techniques, and other advancements that researchers have made in the last decade – but these techniques are exactly what makes a WSN sensitive to interference such as the presence of people. This is a quandary because we often would like to deploy CPS purposely around people, it is the interaction of people-and-system that makes it a CPS; and many of the promises of CPS (smart homes, augmented reality, smart civil infrastructure, disaster mediation, and so on) can be delivered only when CPS are deployed around people daily lives.

One possible (and simple) solution to this quandary, applicable to almost any deployment around people (including deployment with a single base station, and those that would not be able to take advantage of multiple floors and 3D paths) is to install the motes at the ceiling level.

To support future researchers interested in this quandary we mention the following operating characteristics of our network: (1) Relays were operating only with their printed circuit board antennas; (2) The radio operates in the unlicensed 2.4GHz band; (3) All of our relays use the default 26 frequency band; (4) CTP was used for routing. In spot checks that we have performed at the beginning and near the end of the project the electromagnetic radiation at channel 26 was free of wireless interferences. An example to such a spectrum map is presented in Figure 4.5. Following [23] we consider a noise level above -85 dBm noisy enough to cause a Chipcon CC2420 radio to drop packets. In the example provided in the figure we see that channel 26 does not introduce that risk, hence packets (in this case) are not likely to be drooped due to interference from other electronic devices, and the interference must stem from other sources.

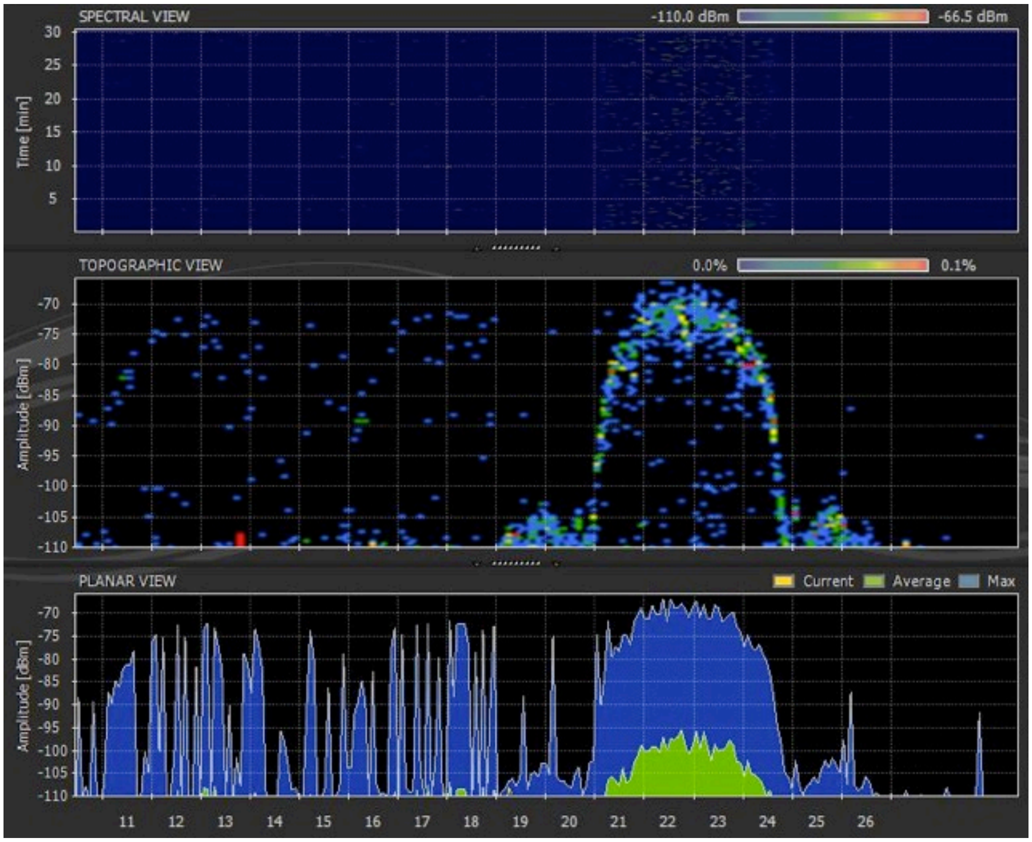


Figure 4.5 Spectrum map accumulated during 30 minutes observation. The X-axis indicates 802.15.4 channels (11 to 26). In the middle and lower figure the Y-axis indicates power. In the middle figure each colored dot (at position x GHz and energy level of y dBm) indicates how often a signal was detected for this ordered pair. The colors represent values in the range shown on the colorbar. Because the area over channel 26 has mostly blue dots (a transmission was detected close to 0.0% in this area) and the power is below -100 dBm (only signals above -85 dBM can interfere with a CC2420 radio) we conclude that channel 26 barley (or not at all) experiance electromagnetic interferences.

Figure 4.5 suggest that frequency hoping in the environment where our system was deployed will not improve the network reliability because channel 26, which we used exclusively in our network, is completely free of interference. In fact channel hoping will only introduce delays, consume energy, and will lower the reliability in such cases because it is bound to hop into noisy channels.

## 4.3 Vital Signs Delivery

The effect of the critical events in the network reliability over the trial's lifetime surely affects the fraction of vital signs actually arriving to a base station from the patient. How bad was it, or how good was it, despite all of these adverse events it what we set to study in this sub-section. Figure 4.6 plots the fraction of each patient's vital signs that were successfully delivered to a base station. Before the July 29 expansion the fluctuating quality of the routing relay infrastructure led to corresponding fluctuations in vital signs delivery. However, after the expansion the reliability improved dramatically. Before the expansion only 25.0% of the patients had at least 70% of their vital signs successfully delivered over the network, compared to 85.4% afterwards. A cursory look at the study coordinator notes on post-expansion patients and base station logs revealed cases of low reliability that can be attributed to extraordinary (not WSN related) factors. For example, one patient was enrolled when base station 52 was off (event (c)). Another was enrolled for only three hours during their hospital discharge, and a detailed examination of her/his database records shows sporadic valid readings, spread among long intervals with no data being received, and intervals of the network successfully delivering finger out of device errors. One probable cause for such a pattern could be that the patient was out-of-range of our deployment span during the discharge process.

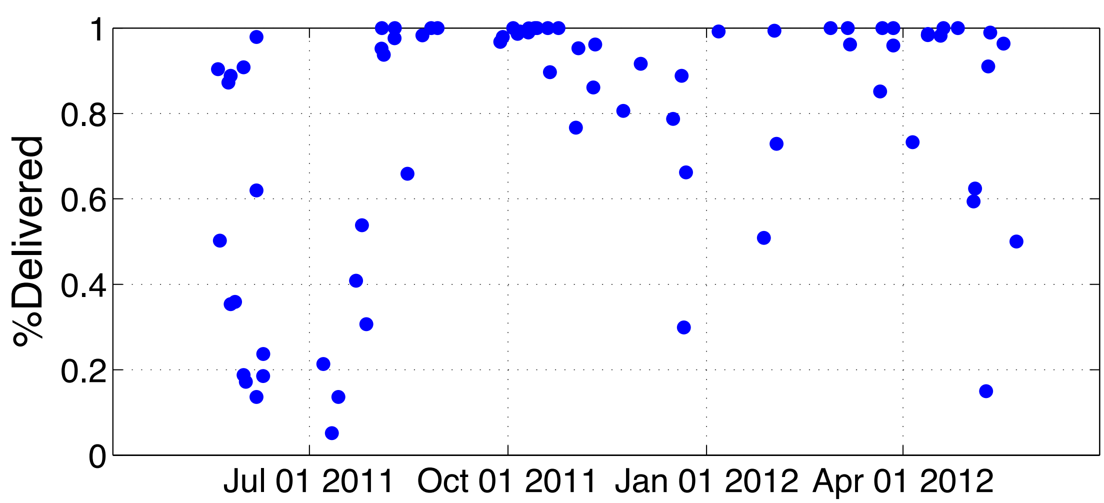


Figure 4.6 The fraction of each patient's vital signs data that is successfully delivered to a base station.

Delivery of at least 70% of the vital signs of 85.4% of patients does suggest that WSN are reliable and can be used for clinical monitoring. First and foremost, by following the insights we offer in Chapter 5, clinical systems can be made much more reliable from the get go. Second, compared to a manual process in which a nurse collects vital signs every few hours we are judging the system's ability to deliver a vital sign every minute. Certainly the average reliability of 85.4% could be improved if we relax this demand and deliver, for example, a vital sign every 5 or 10 minutes.

# Chapter 5

# Lessons Learned

The system has many components such as wireless networking, multiple servers, embedded and regular software, and integration with CDR. We considered all of these pieces within the span of a CPS system, and thus our responsibility. Surprisingly, only few of the challenges we faced stemmed from the many integration points, or as a result of technology in general. While previous research focused on technological solutions, our experiences show that human factors may play at least an equal role in the reliability and usability of a CPS system. In this section we articulate the key challenges and provide suggestions how to avoid them in future deployment. Our discussion revolves around the following key types of challenges: the use of multiple base stations, the need for fault tolerant design, we offer some comments on obtaining correct data for network analysis, we re(discovery) the importance of a well designed user experience, the impact of IT procedures, the effects of the hospital environment, the importance of how things look, and we end with our comments on research administration. Each of the following sections will describe the challenge, the solutions that we took during the on going trial (if it was possible), and recommendations for future research and deployments.

## 5.1 Multiple Base Stations

**Challenge** Multiple base stations can be both a challenge and a blessing. As discussed in section 4.1 the importance of this added infrastructure was quite surprising at first as ward 3 was separated from wards 1 and 2 by a thick ceiling that is full of metal and other hospital building infrastructure. Before the deployment time, it was not clear that ward 3's infrastructure would merge with the existing infrastructure at all, much less significantly reinforce it. We do not attribute this unexpected improvement to an increased network density, since density was comparable in all three wards and adding relays did not improve the reliability.

**Solutions** Previous studies have discussed network density in the context of area coverage, or added relays, but our experience suggests that such coverage alone may not be the only critical factor in heavily populated environments. Our experiences emphasize the importance of leveraging multiple base stations to enhance reliability and the advantages of protocols such as CTP that exploit multiple base stations automatically. Had we deployed the networks in all three wards as separate "sub-networks" instead of allowing them to merge, wards 1 and 2 would have continued to suffer significant data loss, and ward 3 would be less reliable as well. We state two chief design considerations: (1) Multiple base station and their spatial placement (i.e. not just relays spatial distribution) within WSNs are crucial for network connectivity and reliability. (2) Three-dimensional network topologies provide route diversity that enhances reliability in busy buildings, providing new path diversity and thus resiliency against possible interference by humans and other factors.

**Recommendations** Deploy base stations for redundancy. If the physical environment allows, deploy the base stations such that 3D network form a resilient routing infrastructure. While it is possible to form 3D networks with relays alone we believe that there is more path diversity, rendering the network more resilient, when the base stations themselves offer a 3D paths to terminate on a different level than where the destination is. However, when it is not practical use multiple base stations, or to place them at different levels, it is an advantage to deploying the relays on multiple floors. A practical solution that may avoid the problem all together is to install the motes above people's height. As a WSN community we can develop a simple "extension cord apparatus" that powers a mote using a cable extending from the power outlet to the ceiling height, where the mote would reside. This would allow motes to use communication airways that do not go through people, who absorb some of the electromagnetic radiation and thus interfere with the network operation.

## 5.2 Fault Tolerant Systems

### 5.2.1 Integration Points

**Challenge** Some software failures on a base station are particularly disruptive because, in contrast to hardware failures, the base station's attached mote will continue to receive power and act as data sink. Under a hardware failure, relays would begin to re-route their data to another base station under the guidance of a protocol such as CTP. However, under software failures, the base stations effectively converts into a "black hole": receiving power, the base station's motes would still advertise themselves as data sinks, but, but lacking a client to consume the packets the mote would silently discard any incoming data. Hence, even circumstances beyond the control of the data collection system (such as an operating system failures on a base station) can lead to quality degradation. Such monitoring needs to become part of the system's responsibility, despite the fact that the system is not working at the moment. Unfortunately, the nature of the failure means that we cannot estimate the amount of data lost to software failures in our current deployment.

**Recommendations** The system design should incorporate fault detection and recovery mechanisms to deal with both software and hardware failures. We find such design intellectually stimulating in particular because we are asking a system to monitor events when the system itself is not working. For example in the case of the base station's attached mote, it would turn into this "black hole" through which endless data would be lost when the base station software is not present to consume its data. Upon detection of a base station failure the gateway mote attached to the base station should stop broadcasting that it is a viable sink immediately. A specific challenge to the way motes are constructed today is posed because motes cannot query the server to which it is attached. The detection of such a fault can reside with the embedded software on the mote or within the network because those are two components that are properly functioning in this particular example. Focusing on the part of the system that does work in this case yields a fault tolerant solution. The mote software should periodically be receiving status messages from the server and the arrival of such messages indicates that the software is in working condition. However, as soon as such messages stop to arrive, the mote should stop advertising itself as a viable sink.

Similar mechanisms should be designed for each part of the system that may fail. In general we believe that the following design pattern can guide the thinking process.

1. Identify a possible Faulty module of the system, $\mathcal{F}$, which is not functioning under this scenario.
2. Identify modules of the system does do work, $\mathcal{W}_i$, under this this scenario.
3. If a working module $\mathcal{W}_i$ can check a status flag on a non-functioning part $\mathcal{F}$, then the lack of response indicates that $\mathcal{F}$ does not work. Otherwise,
4. Design a simple message exchange protocol such that the lack of arrival of a message from $\mathcal{F}$ to $\mathcal{W}_i$ indicates that $\mathcal{F}$ does not work.

### 5.2.2 Access to Individual Motes

**Challenge** Often during this long study we wished we had direct software access to an individual mote. Either to check its status, know its current location, or query its ID in order to return it to its original location. This is not possible with unidirectional (from leaves (relays) to the root (relay attached to the server)) protocols such as CTP, but it is possible with IPv6 (or similar address based) protocol. We consider the system's ability to query individual components of its infrastructure part of its ability to respond in real-time to events, and part of becoming a fault tolerant system.

**Recommendations** The advantages and disadvantages of address-less vs. address-based protocols is a design decision, per the system particular goals. However, we recommend, for hospital clinical systems (or similar environment in which motes can be moved by people or get lost in a large spatial space) to use address-based protocols such as IPv6. The same recommendation holds if the CPS needs to be fault tolerant because the system needs to be aware of the constituents of its infrastructure. Address-based protocols allow access to each individual device in the network. However, this recommendation poses a new set of challenges for WSN based CPS. These challenges come from the inherent requirement of specifying a destination for each packet. CTP's advantage is that a packet is just being sent into the network and without flooding or broadcasting the packet will arrive to one of the base stations. Flooding and broadcasting are not energy efficient on any network, but this issue is more prominent in the context of WSN and its low energy goals. Additionally, specifying an address for each packet that is sent using an IPv6 or similar protocol, strips away one of the biggest advantages we have discovered in this trial – the base station redundancy. As a reminder (see section 5.1) we hypothesized that this will be a key contributor to network reliability in densely populated areas. It certainly was so in our trial. As the last design consideration we mention that broadcasting in the context of WSN lowers the system scalability by consuming its limited bandwidth. This Thesis is not the place for a new protocol design, nor we have enough information to even suggest that, but we intent to explore protocols that will marry benefits from the two approaches.

## 5.3 Network Analysis

### 5.3.1 Motes Numbering

**Challenge** Each mote in the network had a unique ID that identified its hardware in the cloud. This ID was used to denoted its physical location on the floor layout of the hospital. As noted elsewhere in this paper, in crowded environment motes will tend to disappear and will need to be replaced. Those motes can also reappear in a different room in the hospital, which is the worse case scenario if you are interested in performing a network spatial analysis because this mote ID is associated with its original location.

**Solutions** As soon as we noted that motes could be moved around we trained the hospital staff and our study coordinator about the importance of having a good track over motes locations. It is uncertain to what extent this was effective as we continued to experience motes disappearances until the end of the trial. This should be expected in environments where new people pass by every day.

**Recommendations** Two cases should be considered whenever spatial network analysis is part of the study:

(1) When a disappearing mote is replenished with new mote hardware, the new mote must be assigned the original mote ID.

(2) If a mote ID appears in a new location, it should not be allowed to join the network before its ID is changed.

We recommend separating the TinyOS Active Message ID (AM-ID) from the location ID discussed in this section. AM-ID in TinyOS is determined during compile time and becomes part of the object code. Using this ID as the location ID puts the burden of managing the inventory of location IDs on the researcher, which is error prone and inflexible because the ID is attached to a particular hardware. We recommend that a different set of IDs that can be dynamically assigned will be used for network spatial analysis. The Location-ID (LID) should be managed from a centralized location and an LID consumption mechanism should be developed. The following design objectives should be taken into account:

- Mechanisms to automatically assign LIDs during the initial installation are required. These needs to include assigning LIDs to a large number of motes that have not yet joined the network.
- A researcher that walks into the premises with a mote should not automatically consume a LID as this mote might be unrelated to this deployment.

- A mote should not continue to operate as a relay upon its disconnection from a power outlet. This can be managed by the mote's disconnection alarm (see section 5.5.1).
- Mechanisms to address (1) and (2) above are required.

### 5.3.2 Motes Spatial Location

**Challenge** We found out that in a large study as ours, spanning over large areas, full of other equipment and people, it is not easy to find motes even if you know that they are in a certain room. To be attentive to other hospital priorities and mission we plugged our relay motes into any available power outlet that was not regularly used for medical equipment. We ended up with motes in different locations in each room, which made them hard to find. As a result when we looked for a mote in a room full of other hospital equipment we were often wondering whether it was removed from the room or we are just not looking in the right spot.

**Solutions** In later stages of the trial we implemented the FRBL cardinal direction system described below.

**Recommendations** In future deployment we recommend, based on the available power resources, that a couple of consistent locations would be preferred. Furthermore, each installed mote should be registered with the room number and spot in the room. One such simple scheme that proved beneficial for us is similar to the NESW cardinal direction system. We coined our system the FRBL cardinal direction. In this system we denote on which wall a mote is located relative to a person just entering the room and standing in the doorway. The letter F denotes that the mote is on the wall in Front of the entry door, B for Behind denotes the door's wall (or just Behind the person entering the room), and R/L denote on the Right/Left wall respectively. Of course the known NESW can also be used where N denotes on front wall, E on the right wall, S is the back wall, and W on the left wall. However some people may find this odious because of the probable mismatch with Earth magnetic direction.

## 5.4 End User Indicators

**Challenge** The patient wireless device we used for this trial was developed for a much smaller, previous study, in which it was used successfully. The larger scale of our study coupled with the fact that the patient populations in our study was very frail surfaced a usability design issue that we could not address in the middle of the trial. The existing patient device has no visual or audible indicators. There is no way to know if it is ON or OFF, there is no way to know if it is sensing vital signs (valid or not), and no battery level data is available. The lack of these statuses meant that they could not be sent to the research coordinator or be indicated on the device.

**Recommendations** All of these issues will need to be addressed in future revisions of the patient device. For the author of this paper at least, as an engineer who is always trying to build minimalistic and functional hardware or software that does exactly what it is suppose to do, and no more – this was an excellent lesson in machine-user interaction. For me, putting numerous indicators would not be a question anymore. In a technological age as ours, where storage and to a good extent network bandwidth are not scarce resources anymore, I would even dare to say: the more meaningful indicators there are – the merrier. I draw the same conclusion for building hardware and software that exposes internal indicators (e.g. battery level) and building mechanisms that gives the system (and the users at the UI level) visibility to them.

## 5.5 Impact of IT Procedures

**Challenge** A common challenge for IT in many enterprise networks is ensuring data security and privacy while also giving users enough access to perform their jobs. When moving deployments from a test bed to an enterprise setting it is important to work closely with IT to understand how these policies may affect the deployment. For example, we used laptops as base stations and thus this equipment was treated as end-user equipment. In our case this meant the equipment was preinstalled with a disk image containing mandatory full disk encryption, user-level backups, and a centrally managed software updates policy was enforced. Shortly after the deployment we encountered unanticipated conflicts between these policies and our base stations' normal activities, which were to be performed continuously and in real-time. The end-user level backup software, for example, would mandatory start and block the database server from its real-time work. As another example, after every patch update (e.g. OS, security) that required a reboot the base station would wait for the user to input a password to "unlock" the encrypted hard drive.

**Solutions** We worked with IT to clarify that patient information on these base stations we sanitized according to HIPAA rules. Consequentially the base stations were exempted from full disk encryption and backups[2], and reboots were reduced to once a month for critical patches. We also changed our software to run as an OS service so it will be started automatically after a reboot.

**Recommendations** The integration to the hospital infrastructure and policies are two different objectives, and the integration-design process could be lengthy. These policies would require the system to have additional logic to handle these procedures and this logic may be unrelated to the core system's functions. To incorporate these functions therefore, it is important to start discussions with the organizational IT during the system design. This may lead to selection of alternate equipment such as non-portable or server-class machines, which falls under different policies and often use server-class software that is more suitable for continuous server work. It is also important to test the system on-site with the software functionality that accommodates the hospital policies; as such integration issues would not manifest themselves in standalone WSNs test bed. Some policies should be incorporated into the network

[2] Manual database backups were performed periodically using a utility provided by the database vendor.

design as well, e.g., proactively re-computing routes to avoid machines that are under temporary scheduled maintenance.

## 5.6 Hospital Environment

### 5.6.1 Relay Disconnections

**Challenge** The relay devices and base station were initially deployed without any labels that indicate they were part of a clinical monitoring system. In the beginning of the deployment our relays, base stations, and base stations motes were repeatedly disconnected. The staff indicated that some devices were unplugged and set aside out of curiosity or suspicion. We also believe that some devices have been unplugged by cleaning staff who needed the power outlets since there was no indication that the device was important, nor was it plugged to the hospital medical-class power outlets. It is worth noting that our experience contradicts the common assumption that power is an unlimited resource in indoor WSN deployments because we may leverage available power outlets. Relying on power outlets involves the tradeoff of making the deployment susceptible to disconnections; similar issues were experienced in a residential study [24].

**Solutions** Labeling is crucially important. While we still experience relays disconnections to date, the frequency lowered dramatically after we labeled the relays. As shown in figure 5.1, new relay devices are now deployed with laminated labels. The label is designed to conceal the exposed circuit board, making them appear less suspicious. Similar labels have been attached to the base stations' lids to discourage disconnections or personal use and prevent disconnections of their attached motes.


Figure 5.1 The relay device before and after labeling.

### 5.6.2 Base Stations Disconnections

**Challenge** During the trial, base stations were disconnected from the power, network, or attached motes. These base stations act as bridges between the WSN and the hospital's CDR database and such disconnects degraded the quality of the system. Base station disconnections were particularly surprising: in contrast to the relay nodes, which were deployed in patient rooms and hallways, the base stations were in rooms only accessible to hospital staff. We soon discovered that location sends a message.

**Solutions** We used laminated labels for the base stations, their attached motes, and power cables; as well as relays.

**Recommendations** With these experiences in mind we have a number of recommendations: (1) Equipment should not only be able to perform its task but also professionally display its purpose (e.g., base stations should be server-class machines, as much for appearance, as for the robust hardware). (2) Equipment should be installed in appropriate locations that discourage casual tampering, personal use, or disconnections. (3) Labels should clearly indicate that the base station is medical equipment. (4) Labels should warn against disconnection on cables near susceptible outlets or ports (when someone is under the desk looking for an Ethernet port or a power outlet they do not know which cable is medically important).

### 5.6.3 Disconnection Alarm

**Challenge** We plugged our relay motes into power outlets but they did not always stay there. Operating under the assumption that our motes will have limitless power (being plugged into an outlet) we did not equip them with batteries. In a busy hospital it can be expected that devices will be unplugged and reappear in a different location or not plugged back at all.

**Solutions** It was not practical for us to use a disconnection alarm because it would require someone to be present in the hospital at all times to attend to it and we were budgeted only for a part time study coordinator.

**Recommendations** In future studies we recommend to equip the motes with batteries despite the fact they will be plugged into a power outlet. A disconnection mechanism can then send an alert, immediately upon disconnect, to a monitoring module of the system. We also believe that sounding an audio alert upon disconnection may encourage the person to return the mote to its place immediately. The battery will also allow the mote to stay in communication with the monitoring module and indicate its current location. The signal strength relevant to other motes, or even its path to the base station can be used as location indicators.

### 5.6.4 Professional Medical Standards

**Challenge** There are a number of topics that fall under this category. They are all discussed in the recommendation section.

**Recommendations** We already discussed using server-class machines as medical equipment in section 5.6.2. In fact, with our desire to run everything on standard hardware almost every project that uses regular computers for specialty-use (e.g. as medical equipment) runs the risk of being treated as a computer for surfing the Internet. During research, when it does not make sense to externally brand your equipment, one good course of action is to use machines that look a little bit different. Standardize on the same equipment in all locations, install them in "medically reserved spots", and clearly label the equipment. This will brand your equipment as close as you can to a medical equipment during the research period and may prevent tampering with its operation.

Health risks are a hazard that hospital staff is aware of, and is trained to deal with. However this may be new ground for computer scientists and engineers. Get trained before the deployment starts. Many, if not all, hospitals have a department responsible for sterilization procedures. Get in touch with them and be consulted on proper sterilization of the equipment between patients or before equipment is returned to the lab for maintenance (you may want to note that in the lab journal). This may be particularly important if other students or lab mates will use the equipment after the trial has ended.

If your equipment is not disposable then after sterilization we recommend it would be sealed and packaged as appropriate. This makes the device (and the project as a whole) appear more professional. We also believe it will make the volunteers feel more comfortable because this is what they see when commercial equipment or utensils are used during their treatment.

Finally we recommend that you will get a lab coat, it usefully reduces the risk of carrying pathogens back to your car, office, or home; and it makes you part of the professional hospital team.

# 5.7 Research Administration

## 5.7.1 Lab Journal

**Challenge** There are many events that transpire in a study that spans a long duration or large grounds. For example, patients may decide to stop wearing the device on their own or a base station may be shutdown by someone. To correlate events with the system or network performance we sifted through emails, but this is very inconvenient, time consuming, and inaccurate method. It is almost guaranteed that you will not find a record of everything that occurred during the trial.

**Solutions** To the extent that we could, we archived every research related email in personal folders. Some people switched to gmail for its superior search.

**Recommendations** A properly managed lab journal (paper based or electronic) will help you get closer to the truth when it comes to events in the project life.

# Chapter 6

# Related Work

Industry is offering wireless telemetry solutions such as the GE Healthcare's ApexPro telemetry system [25], Sotera WIRELESS's ViSi Mobile System [26], and Philips's IntelliVue cableless solution [27]. Each of those technologies (and our own) presents different strengths and limitations. The commercial wireless solutions require an installation of dedicated infrastructure such as wired access points (i.e. a WiFi AP requires a wired Ethernet connection) which are labor intensive to install and costly. In contrast, WSN technology employs low-power wireless mesh networking that can be easily deployed without any fixed infrastructure (other than power outlets). The cost-effectiveness and ease of deployment makes WSNs particularly attractive for resource constrained clinical settings such as field hospitals, rural areas, and developing countries. Lastly, these wireless technologies require more power than WSN, which is based on low-power wireless standards such as IEEE 802.15.4 [12]; hence portable devices based on low-power WSNs may be deployed with slimmer batteries or may require less frequent charging.

Previous studies have found great promise in WSNs for emerging medical applications such as emergency care in disaster areas [4, 28, 29], assisted living and residential monitoring [6], and early detection of clinical deterioration [3, 4]. In contrast to the study described in this paper, these prior studies were deployed as standalone WSNs and were not integrated with the legacy hospital systems. Moreover, these deployments were limited to a single unit or a relatively small area. As we have seen in our study larger spans of WSN call for additional design considerations.

# Chapter 7

# Conclusions

Wireless networks have shown promise for real-time clinical monitoring in hospitals. In this work we evaluate the question whether a subset of wireless network, namely low-power wireless, can be used to build a reliable clinical monitoring cyber physical system. We focus on low-power networks because of advantages such as easy deployment and energy consumption. This paper presents the lessons and insights learned from a 14-month deployment of a large-scale WSN for vital sign monitoring in a major hospital. Salient features of our system are scale (spanning four hospital floors) and integration with existing IT infrastructure and CDR systems in a hospital. We present important and sometimes surprising findings that were not reported in previous WSN deployments in hospitals and other environments. (1) Base station placement is crucial for reliable data collection and allowing subnetworks to communicate with each other can effectively increase the overall system's reliability and render it more resilient to hospital events. (2) Network reliability can be increased by designing for 3D network topologies. (3) Fault tolerant design is necessary to progress towards production-wordy CPS deployments. (4) Designing meaningful end-user indicators can make a positive impact on a CPS. (5) It is crucial to deal with equipment disconnections even in a well-controlled hospital environment, and to establish protocols that recover and proactively discourage these disconnections.

Our results lead to guidance and best practices of deploying large-scale WSNs in hospital environments and other large-scale WSNs.

# References


[1] G. Hackmann, M. Chen, O. Chipara, C. Lu, Y. Chen, M. Kollef, and T. Bailey. *Toward a two-tier clinical warning system for hospitalized patients*. AMIA, 2011.

[2] K. Zeitz and H. McCutcheon. *Policies that drive the nursing practice of postoperative observations*. International journal of nursing studies, 39(8):831–9, Oct 2002.

[3] O. Chipara, C. Lu, T. C. Bailey, and G.-C. Roman. *Reliable clinical monitoring using wireless sensor networks: Experience in a step-down hospital unit*. In SenSys, 2010.

[4] J. Ko, J. H. Lim, Y. Chen, R. Musvaloiu-E, A. Terzis, G. M. Masson, T. Gao, W. Destler, L. Selavo, and R. P. Dutton. *Medisn: Medical emergency detection in sensor networks*. ACM Transactions on Embedded Computing Systems, 10(1):11:1–11:29, Aug 2010.

[5] K. Lorincz, D. Malan, T. Fulford-Jones, A. Nawoj, A. Clavel, V. Shnayder, G. Mainland, M. Welsh, and S. Moulton. *Sensor networks for emergency response: challenges and opportunities*. IEEE Pervasive Computing, 3(4):16–23, 2004.

[6] A. Wood, G. Virone, T. Doan, Q. Cao, L. Selavo, Y. Wu, L. Fang, Z. He, S. Lin, and J. Stankovic. *Alarm-net: Wireless sensor networks for assisted-living and residential monitoring*. Technical Report CS-2006-11, University of Virginia, Dec 2006.

[7] Y. Mao, Y. Chen, G. Hackmann, M. Chen, C. Lu, M. Kollef, and T. Bailey. *Early deterioration warning for hospitalized patients by mining clinical data*. International Journal of Knowledge Discovery in Bioinformatics, 2(3):1–20, 2012.

[8] Memsic. (2013, April) *Telosb mote platform*. http://www.memsic.com/wireless-sensor-networks/TPR2420

[9] J. Hill, R. Szewczyk, A. Woo, S. Hollar, D. Culler, and K. Pister. *System architecture directions for networked sensors*. SIGPLAN Notices, 35(11):93–104, 2000.

[10] Covidien. (2013, April) *Nellcor pulse oximetry sensors*. http://www.covidien.com/rms/pages.aspx?page=OurProducts/Oximetry/Sensors

[11] Smith Medical. (2013, April) *OxiLink*. http://www.smiths-medicaloem.com/catalog/monitoring-technology/patient-monitoring/oximetry/oxilink.html

[12] Texas Instruments. (2013, April). *Single-Chip 2.4 GHz IEEE 802.15.4 Compliant and ZigBee™ Ready RF Transceiver*. http://www.ti.com/product/cc2420

[13] O. Chipara, C. Brooks, S. Bhattacharya, C. Lu, R. D. Chamberlain, G. C. Roman, and T. C. Bailey. *Reliable real-time clinical monitoring using sensor network technology*. In AMIA, 2009.

[14] IEEE. (2013, April). *IEEE 802.15 WPAN Task Group 4 (TG4)*. http://www.ieee802.org/15/pub/TG4.html

[15] O. Gnawali, R. Fonseca, K. Jamieson, D. Moss, and P. Levis. *Collection tree protocol*. In SenSys'09.

[16] The Jython project. (2013, April). http://www.jython.org/index.html

[17] Tanuki software. (2013, April). Jave service wrapper. http://www.tanukisoftware.com/en/index.php

[18] PostgreSQL. (2013, April). http://www.postgresql.org/

[19] Health Level Seven International. (2013, April). http://www.hl7.org/

[20] HAPI. (2013, April). http://hl7api.sourceforge.net/

[21] WebSphere MQ. (2013, April). *A flexible system for efficient transport of messages and data*. http://www-03.ibm.com/software/products/us/en/wmq/

[22] Java Native Access (JNA). (2013, April). http://twall.github.io/jna/3.5.2/javadoc/overview-summary.html#overview_description

[23] M. Sha, G. Hackmann and C. Lu. *Real-world Empirical Studies on Multi-Channel Reliability and Spectrum Usage for Home-Area Sensor Networks*. IEEE Transactions on Network and Service Management, 10(1): 56-69, March 2013. (Extended version of the IWQoS'11 paper).

[24] T. Hnat, V. Srinivasan, J. Lu, T. Sookoor, R. Dawson, J. Stankovic, and K. Whitehouse. *The hitchhiker's guide to successful residential sensing deployments*. In SenSys, 2011.

[25] GE Healthcare. (2013, April). *GE ApexPro Telemetry System*. http://www3.gehealthcare.com/en/Products/Categories/Patient_Monitoring/Wireless_Networks

[26] Sotera WIRELESS. (2013, April). *ViSi Mobile System*. http://www.visimobile.com/

[27] Philips. (2013, April). *Monitoring Wireless Solutions.* http://www.healthcare.philips.com/us_en/products/patient_monitoring/products/wireless_monitoring_solutions/index.wpd#

[28] T. Gao, et-al. *The advanced health and disaster aid network: a light-weight wireless medical system for triage.* IEEE Transactions on Biomedical Circuits and Systems, 1(3):203–216, Sept. 2007.

[29] D. Malan, T. Fulford-jones, M. Welsh, and S. Moulton. *Codeblue: An ad hoc sensor network infrastructure for emergency medical care.* In International Workshop on Wearable and Implantable Body Sensor Networks, 2004.

[30] IEEE. (2013, April). *IEEE 802.11 Wireless Local Area Networks (WLAN).* http://grouper.ieee.org/groups/802/11/

[31] John Bicket, Daniel Aguayo, Sanjot Biswas, Rober Morris. *Architecture and Evaluation of an Unplanned 802.11b Mesh Network.* In MobiCom 2005.

[32] Qualcomm Life. (2013, April). *The 2net Platform.* http://www.qualcommlife.com/wireless-medical-information-technology

[33] White Space. (2013, April). http://www.weightless.org/