Experiences with an End-To-End Wireless Clinical Monitoring System

Rahav Dor[1], Gregory Hackmann[1], Zhicheng Yang[1], Chenyang Lu[1], Yixin Chen[1], Marin Kollef[2], Thomas Bailey[2]
[1]Department of Computer Science and Engineering [2]School of Medicine
Washington University in St. Louis
{rahav.dor, ghackmann, yangzhicheng, lu, ychen25, kollefm, tbailey}@wustl.edu

ABSTRACT
Wireless sensor networks can play an important role in improving patient care by collecting continuous vital signs for clinical decision support. This paper presents the architecture of, and our experiences with, a large-scale wireless clinical monitoring system. Our system encompasses portable wireless pulse oximeters, a wireless relay network spanning multiple hospital floors, and integration into the hospital Electronic Medical Record (EMR) databases. 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. 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.

Categories and Subject Descriptors
C.2.4 [**Computer-Communication Networks**]: Distributed Systems; J.3 [**Computer Applications**]: Life and Medical Sciences—*Medical information systems*

General Terms
Design, Experimentation, Human Factors, Reliability, Verification

Keywords
Wireless Clinical Monitoring, Wireless Sensor Networks

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 [5]. Wireless sensor network (WSN) technology can enable real-time detection of clinical deterioration by collecting continuous vital signs from patients. Compared to commercial telemetry systems, WSNs based on low-power technology have advantages in both power efficiency and cost-effective deployment without fixed infrastructure [2].

While earlier deployment of WSNs in hospitals have shown promise [2, 7, 8, 12], they are not sufficient for clinical detection because their *standalone* WSNs are not integrated with the hospital's Electronic Medical Records (EMR) 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 unit, while large-scale and long-term deployments spanning multiple clinical units are needed to demonstrate the feasibility of WSN technology in hospitals.

This paper presents the architecture and deployment of an *end-to-end* clinical monitoring system that integrates portable wireless pulse oximeters for patients, wireless relay networks spanning multiple hospital floors and wards, and departmental servers communicating with the hospital's Electronic Medical Record (EMR) database. This study is distinguished from the above trials based both on the EMR integration and the study's scale. 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 EMR in real time, where it resides along with an historical corpus of clinical data. This extended EMR data can be used to detect clinical deterioration using data mining techniques [5].At the time of this writing, our 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 existing hospital EMR systems.

Although our deployment's EMR integration, spatial scope, and duration may appear to be subtle differences from prior studies, they 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 both a challenge and a blessing, increasing the network reliability and resilience to hospital events. (2) As part of the planning process, human factors such as the behavior of patients, clinicians, hospital staff, and visitors must be addressed *before* the deployment

Wireless Health '12, October 23-25, 2012, San Diego, USA

starts. (3) 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 around people's daily lives.

The rest of this paper is organized as follows. Section 2 describes the study that provides the context of the clinical monitoring system. Section 3 discusses the system architecture. Section 4 presents our findings about the system's reliability during the clinical trial, in terms of both network performance and reliability of vital signs delivery. Section 5 articulates the lessons we learned and how the overall performance of such a system can be improved by following a few simple guidelines. We end this paper with a short review of similar efforts in the community and our conclusions.

2. CLINICAL TRIAL OVERVIEW

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

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 timely attention to their patients. In order to achieve these goals we envision a novel two-tier system for clinical early warning [5]. 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 which 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.

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 [5], followed by an in-depth study of data mining algorithms for traditional clinical data presented in [10]. The success of such systems depends on the availability of a reliable real-time feed of patients' vital signs integrated into existing EMR systems. This paper presents a clinical trial of a real-time clinical monitoring system that collects vital sign data from wireless sensors, transport, and feeds it to EMR. Once in EMR the real-time vital signs along with the traditional clinical data of multiple patients can be used for detecting imminent clinical deterioration. This phase of the deployment focuses on gathering sufficient data to determine if this can be done using WSNs. Data collected during this trial was not used to predict patient outcomes or manage patient care. In the following sub-section we provide some details on the participant sample and anecdotal field observations.

Status	Count	Percent
Consented	75	21.1
Consented, but no data collected	12	3.4
Refused consent	42	11.8
Contact isolation	131	36.9
Not available or responsive	67	18.9
Do No Resuscitate (DNR)	15	4.2
Not appropriate per nursing	10	2.8
Dementia	3	0.9

Table 1: Patient consent results

Race	Count
Black	30
Caucasian	40
Other	5

Gender	Count
Female	51
Male	24

Table 2: Demographics of consented patients

2.1 Patient enrollment

The tier-1 data mining algorithm runs on EMR database to identifies patient at risk, and alerts the study coordinator. The coordinator then tried to consent the patient to participate in the clinical trial; Table 1 presents the results of this enrollment process. Of the total of 355 patients tier-1 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), and only 42 refused consent. The remaining patients did not have the opportunity to consent due to their medical condition. An additional 173 patients were excluded from the table because they were discharged before enrollment.

Table 2 presents the volunteers demographics. During the trial 2 patients expired and 4 were transferred to the ICU. The mean age of participants was 58 years.

2.2 Feedback from the field

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.

3. SYSTEM ARCHITECTURE

In this section, we discuss the architecture of the clinical monitoring system. 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 them 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 which carries the vital sign data from the patient devices to 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 EMR database. They also provide network monitoring and logging services that have been invaluable to system maintenance during our clinical trial.

The patient device and wireless relay networks reuse the same design from an earlier clinical trial detailed in [2], while the base station software has been extended substantially with new services to support the EMR integration and network monitoring services. The key distinctions between this study and the earlier study reported in [2] are two-fold: (1) the integration with the hospital IT infrastructure and EMR compared to standalone WSN. (2) a large-scale, multi-floor and multi-unit deployment compared to single-unit deployment. Each of these characteristics lead to important new challenges and insights which we report in this paper.

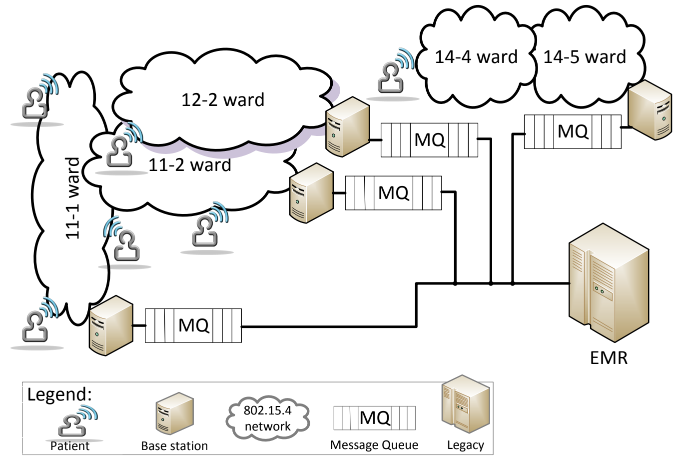

Figure 1: A high-level view of the network architecture. The figure shows the ward floor and number, e.g. 11-1 is ward 1 on the 11th floor.

3.1 Patient device

Each patient device integrates a TelosB mote with a Nellcor pulse oximeter sensor attached to the patient's finger to collect pulse rate and blood oxygenation level (S_pO_2). After receiving a sensor reading the TelosB routes the data to the relay infrastructure discussed in the following subsection, using its onboard CC2420 [11] radio. The patient locate relays using the Dynamic Relay Association Protocol (DRAP) [1], a one-hop protocol specifically designed to locate a larger routing infrastructure under mobility.

3.2 Relays network infrastructure

WSN network coverage is achieved by plugging TelosB nodes into electrical outlets in patients' rooms, offices, and elsewhere when it is 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 [11]. Figure 1 presents a high-level view of the resulting network topology. The Collection Tree Protocol (CTP) [4] is employed on these 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 which uses periodic beacon messages to maintain the routing infrastructure, autonomously changing the links used to route packets as their reliability varies over time.

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 2 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 (light blue) and one base station (dark purple). Black arrows indicate the links used to form the network's routing trees.

We expected that the respective wards would generally form their own isolated routing infrastructures 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 an other ward. However, as shown in figure 2, the relays routinely formed links across wards. Moreover, because the 12th floor deployment was directly above the 11th floor, routing trees often crossed through floors. (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). 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 floor. We discuss both observations later in the paper.

3.3 Base stations

Each base station consists of a laptop with a USB-attached TelosB mote, and acts as an interface between the 802.15.4 network and the hospital's wired Ethernet infrastructure. Each base station is deployed with a *base station manager* software with the core functionality depicted in figure 3. As vital signs produced by patients arrive from the relays to the mote attached to one of the base stations, the vital sign data are sent to the software using the mote's RS232 interface.

The vital signs are then formatted according to the interna-

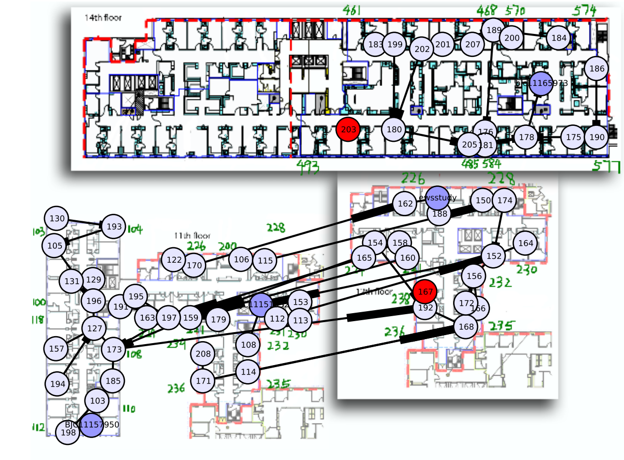

Figure 2: Relays infrastructure across five hospital wards. Light purple circles denote the *relays*. Darker purple denote *base stations*. Red shows relays that are currently *disconnected* from the base stations. The bottom left side of the image shows the network on 11th floor and its two base stations (50 and 52). The bottom right shows the network on the 12th floor and base station 60. The top shows the 14th floor and base station 73.

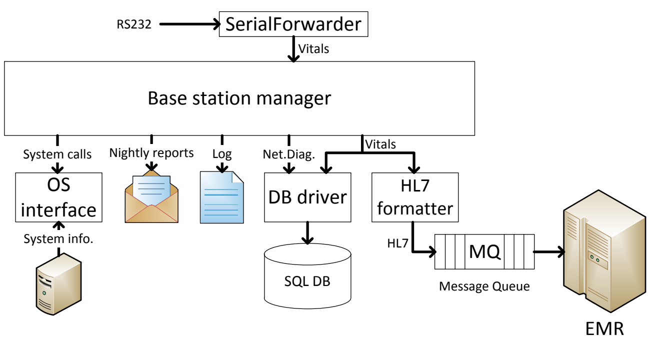

Figure 3: Data flow through the data collection system

tional Health Level Seven (HL7) standard with a converter built on top the HAPI library. This functionality is depicted by the module named *HL7 Formatter*. The HL7-formatted vital sign data are inserted to a *Message Queue* (*MQ*) that reliably transports them to the hospital EMR, over the hospital's wired Ethernet network. This particular deployment used the IBM WebSphere MQ library for consistency with existing hospital IT systems.

An important feature of our base station software is a suite of network monitoring services which have proven invaluable for the maintenance. For diagnostic and debugging purposes, each base station is equipped with a database server (in this implementation PostgreSQL) which is used to locally log all messages received from the relay network. This log includes the relays' diagnostic beacons, which are not forwarded to the hospital EMR, as well as a local copy of all collected vital sign data. Extraordinary events are also logged in a timestamped text file for easy access.

To track the base stations' health each base station produces and emails a *nightly report* every 24 hours (this is 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. 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.

To help us understand the network performance we developed a utility that queries each base station's diagnostic logs to map the network topology during the last 15 minutes. Figure 2 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.

4. NETWORK RELIABILITY

In this section, we analyze our deployment's network 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 deployment's ability to collect real-time vital signs using a WSN. Due to its location the 14th floor formed its own network separate from the rest of the network. As this single-unit network is similar to our previous deployment studied in [2], henceforth our analysis focuses on the challenges posed by the larger, multi-ward deployment on the 11th and 12th 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. We define the PDR threshold to designate relays as *connected* or *disconnected* each day: to be *connected*, a relay must have successfully delivered 90% of its beacons to the base stations that day. Figure 4 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; highlighted regions mark significant events during the deployment that we discuss below. We note that Figure 4 begins at May 24, 2011: due to factors described in Section 5.3 the diagnostic logs needed to produce this graph are partially incomplete between the initial deployment and May 23.

Starting in April 2011, the network was deployed with 11 relays and two base station (50 and 52) at two corresponding

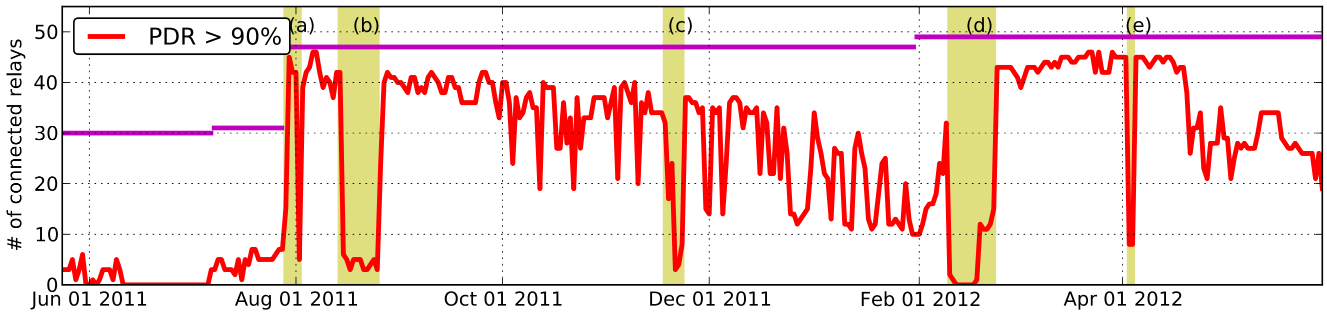

Figure 4: The network reliability from May 2011 to May 2012. The purple line marks the number of deployed relays. The red plot represents the number of relays connected to the base stations. Significant events are shaded: (a) The deployment was expanded to ward 3 on July 29. (b) Between August 14 to 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.

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, the deployment was expanded, adding ward 3 exactly one floor above ward 2. Ward 3 was furnished with its own base station (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.

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 floors, routes that to some extent avoided links through corridors which are usually packed 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: 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 $\geq$ 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

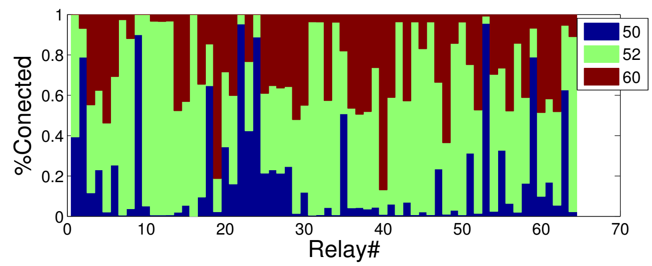

Figure 5: The percent of time relays were connected to a base station (50, 52, and 60). We see that almost every relay in the deployment divided its time among multiple base stations.

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 5 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 are offered in the next section.

4.2 Vital signs delivery

The effects of these critical events over the trial's lifetime may be seen in the fraction of vital signs actually arriving to any base station from each patient. Figure 6 plots this

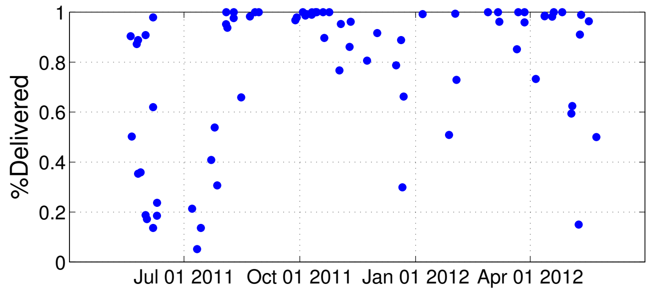

Figure 6: The fraction of each patient's vital signs data that are successfully delivered to a base station.

data. Before the July 29 expansion the fluctuating quality of the routing infrastructure lead 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 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 shows sporadic valid readings spread among long intervals with no data being received (possibly indicating the patient was out-of-range of the network) and intervals of the network successfully delivering *finger out of device* errors.

5. LESSONS LEARNED

The system has many components such as wireless networking, highly available servers and software, and integration with EMR. We considered all of these pieces within the span of responsibility of our system. 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 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 software, the impact of IT procedures, and the effects of an hospital environment. Each of the following sections will describe the challenge, the solutions that we took during the on going trial, and recommendations for future 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. 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.

Solutions: Previous studies have discussed network density in the context of area coverage, but our experience suggests that coverage alone may not be the only critical factor in heavily-trafficked environments. Our experiences emphasize the importance of leveraging multiple base station 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 "subnetworks" 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 (not just relays) within WSNs are crucial for network connectivity and reliability. (2) Three-dimensional network topologies provide route diversity that enhance reliability in a busy buildings, with possible interference from humans and other factors.

Recommendations: Deploy base stations for redundancy. If the physical environment allows, deploy 3D network to form a resilient routing infrastructure.

5.2 Fault tolerant software

Challenge: Some software failures on a base station are particularly disruptive because, in contrast to hardware failures, the base stations' attached motes will continue to receive power and act as data sinks. Under a hardware failure, relays would begin to re-route their data to another base station under the guidance of the CTP protocol. However, under software failures, the base stations effectively convert into a "black hole": the base stations' motes would still advertise themselves as data sinks, but, lacking a client to consume the collected packets would silently discard any incoming data. Hence, even circumstances beyond the control of the data collection system (such as operating system failures on a base station) can lead to quality degradation and *should become* part of the system's responsibility. Unfortunately, the nature of the failure means that we cannot estimate the amount of data lost to software failures in our current deployment. But we can design to address it in future deployments by incorporating fault-tolerant mechanisms in the base station software design.

Recommendations: Base station software should adopt fault detection and recovery mechanisms to deal with both software and hardware failures. 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.

5.3 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 testbed to an enterprise setting, it is important to work closely with IT to understand how these policies may affect the deployment. 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 centrally-managed software updates policy was enforced. Shortly after deployment, we encountered unanticipated conflicts between these policies and our base stations' normal activities, which were to be performed continuously and in real-time. For example, the end-user level backup software would 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 was sanitized according to HIPAA rules. The base stations were exempted from full-disk encryption and backups [1], 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 will affect the system operational logic; hence it is important start these discussions 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 testbed. Some policies should be incorporated into the network design as well, e.g., proactively recomputing routes to avoid machines that are under temporary scheduled maintenance.

5.4 Hospital environment

5.4.1 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 trade-off of making the deployment susceptible to disconnections; similar issues were experienced in a residential study [6].

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 7, 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.

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

Recommendations: In addition to labeling, our recommendation (which would be a beneficial to indoor outlet-powered WSNs in general) is to encase the relays in a small plastic box that is plug-able to a power receptacles, and makes power receptacle available on its faceplate. Such an apparatus will not force cleaning crews to unplug the motes when they need a power outlet and will give a professional appearance that discourages unplugging.

Figure 7: The relay device before and after labeling.

5.4.2 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 EMR 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 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 or disconnections. (3) Use labels clearly indicating that the base station a medical equipment. (4) Use labels warning against disconnection on all equipment and near outlets (when someone is under the desk looking for an Ethernet port or a power outlet they do not know which cable is medically important).

6. RELATED WORK

Industry is offering wireless telemetry solutions such as the GE Healthcare's ApexPro telemetry, Sotera Wireless's ViSi sensor, and Philips's IntelliVue cableless solution. Each of those technologies (and our own) present different strengths and limitations. These commercial solutions require an installation of dedicated infrastructure such as wired access points 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 readily available 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 WSNs, which is based on low-power wireless standards such as IEEE 802.15.4 [2]; 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 [3, 9, 7], assisted living and residential monitoring [12], and early detection of clinical deterioration [7, 2]. 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.

7. CONCLUSION

Wireless sensor networks have shown promise for real-time clinical monitoring in hospitals. 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 EMR 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) While previous deployment studies usually focused on WSNs as a standalone network, our system experienced a significant number of failures caused by interactions with enterprise IT methods and regulations. (2) Base station placements is crucial for reliable data collection and multiple subnetworks can effectively increase reliability. (3) Network reliability can be increased by 3D topologies. (4) 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.

Acknowledgments

We would like to acknowledge the study coordinators Emily Kuo and Pam Kemp for their tireless and caring work with patients, Kevin Heard for superior programming on EMR and for providing us with invaluable data, and Kelly Faulkner for great analysis and keeping us honest to our tasks. Special thanks we send to all the nurses who supported the project.

This publication 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. Its contents are solely the responsibility of the authors. Additional funding was provided by the Barnes-Jewish Hospital Foundation and NSF through grants CNS-1035773 (CPS), CNS-1144552 (NeTS), and CNS-0708460 (CRI).

8. REFERENCES

[1] 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.

[2] 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.

[3] T. Gao, T. Massey, L. Selavo, D. Crawford, B. rong Chen, K. Lorincz, V. Shnayder, L. Hauenstein, F. Dabiri, J. Jeng, A. Chanmugam, D. White, M. Sarrafzadeh, and M. Welsh. 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.

[4] O. Gnawali, R. Fonseca, K. Jamieson, D. Moss, and P. Levis. Collection tree protocol. In *SenSys*, 2009.

[5] 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.

[6] 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.

[7] 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 Trans. Embed. Comput. Syst.*, 10(1):11:1–11:29, Aug 2010.

[8] 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.

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

[10] 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.

[11] Texas Instruments. *2.4 GHz IEEE 802.15.4 / ZigBee-ready RF Transceiver.*

[12] 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.