

Reliable Real-time Clinical Monitoring Using Sensor Network Technology

Octav Chipara, Christopher Brooks, Sangeeta Bhattacharya, Chenyang Lu,
Roger D. Chamberlain, Gruia-Catalin Roman, Thomas C. Bailey
Washington University in St. Louis, St. Louis, MO

Abstract

We propose wireless sensor networks composed of nodes using low-power 802.15.4 radios as an enabling technology for patient monitoring in general hospital wards. A key challenge for such applications is to reliably deliver sensor data from mobile patients. We propose a monitoring system with two types of nodes: patient nodes equipped with wireless pulse oximeters and relays nodes used to route data to a base station. A reliability analysis of data collection from mobile users shows that mobility leads to packet losses exceeding 30%. The majority of packet losses occur between the mobile subjects and the first-hop relays. Based on this insight we developed the Dynamic Relay Association Protocol (DRAP), an effective mechanism for discovering the right relays for patient nodes. DRAP enables highly reliable data collection from mobile subjects. Empirical evaluation showed that DRAP delivered at least 96% of data from multiple users. Our results demonstrate the feasibility of wireless sensor networks for real-time clinical monitoring.

Introduction

Early detection of clinical deterioration in patients is a key factor in reducing mortality rates and length of stay in hospitals. The prevalence of clinical deterioration resulting in cardiopulmonary or respiratory arrests for patients admitted to hospitals is between 4-17%¹. An automated scoring system that uses vital signs and other commonly recorded clinical data has been shown to identify patients at risk of clinical deterioration². Similarly, early detection and treatment of patients with sepsis, an infection induced syndrome resulting in systemic inflammatory response complicated by dysfunction of at least one organ, resulted in statistically significant lower mortality rates³.

While the benefits of early detection are real, the sensitivity and accuracy of such systems hinges upon having timely clinical data¹. This is not a major problem for patients in Intensive Care Units (ICUs) since their vital signs are monitored by wired electronic monitoring systems. However, due to the significant cost of monitoring systems, the vital signs of patients in general wards are most often collected manually, and are obtained at relatively long intervals. For example, even in postoperative care, nurses measure vital signs only 10 times in the first

24 hours following an operation⁴. This may lead to either missing, or prolonged delays in the detection of clinical deterioration. Therefore, *a real-time and reliable clinical monitoring system for patients in general wards is critical for effective early detection of clinical deterioration*. The integration of real-time clinical data automatically collected from patients into electronic record systems will enable advanced clinical applications. For example, an electronic scoring system can automatically generate alerts to a rapid response team in the event of patient deterioration, thereby reducing the delay of response from hours to minutes.

Emerging wireless technologies hold the promise to meet the challenge of real-time clinical monitoring. Commercial medical device vendors (e.g., Phillips, GE, Cisco) are moving toward solutions that use IEEE 802.11 in their wireless patient monitoring systems. These systems require the deployment of wireless routers connected to a wired infrastructure. As a result, the equipment and deployment costs of such systems are high, limiting their adoption in specialized departments such as cardiology. In contrast, our goal is to develop an inexpensive real-time wireless clinical monitoring system for patients in general hospital wards.

The system we developed employs emerging wireless sensor network (WSN) technology based on IEEE 802.15.4 standard. Such a solution is attractive for patient monitoring for the following reasons. First, in contrast to 802.11 radios, 802.15.4 radios may be integrated with inexpensive microcontrollers and hardware platforms. In fact, the retail price of 802.15.4 radios is approaching \$1, and single-package radio, processor, and memory are about \$3¹³. Second, we reduce the *deployment cost* of the system through mesh networking technologies, which require minimal infrastructure. In contrast to current systems in which a patient must be within the radio range of an access point connected to a wired local area network, our system may route data to a base station over multiple hops. Thus, we remove the expensive task of wiring access points. Third, 802.15.4 radios consume significantly less energy (e.g., in idle state a typical 802.15.4 draws 17.4 mA, compared to 200 mA for 802.11). The *small form factor and weight* coupled with proper *power management* techniques that allow operating the device for the duration of hospitalization without




Figure 1. Circles denote relay nodes that are part of testbed. Diamonds indicate relay nodes plugged into wall outlets. A square denotes the base station. Positions A and B are used in the mobility study.

replacing batteries thereby reducing the inconvenience of wearing a monitoring device. Therefore, even in hospitals where 802.11 infrastructure has already been deployed, 802.15.4 technology still delivers important advantages in terms of power consumption and patient convenience.

For a wireless clinical monitoring system to be widely used in hospitals, it is essential to *reliably* collect data from patients. We identified two key challenges in developing such a system: (1) In contrast to 802.11 radios, low-power 802.15.4 radios are unreliable due to dynamic channel conditions as well as susceptible to interference⁹. (2) Unlike patients in ICUs, patients in the general hospital wards may be ambulatory. Patient mobility may lead to significant packet losses. This work focuses on the challenge of reliable clinical monitoring of mobile subjects. Specifically, the contributions of this paper are: (1) Through systematic empirical studies on a prototype system we show that the Collection Tree Protocol (CTP)⁹, the standard data collection protocol on WSNs, performed poorly under mobility. We identify the root cause of these problems to be the first hop; that is, the changing connectivity between the mobile node and the relay nodes. (2) We propose a robust data collection mechanism called *Dynamic Relay Association Protocol* (DRAP) which allows a patient node to dynamically associate with a reliable relay node as the user moves in the environment. (3) Empirical results show that DRAP achieves at least 96% of the sensor readings to the base station and has a radio duty-cycle (the fraction of time the radio is turned on) of 2.2%, thus conserving power.

Related Work

The mobile ad hoc network community has proposed numerous routing protocols with support for mobile entities¹⁰. It is common for routing protocols to introduce significant overhead by reconstructing a portion or an entire route in response to mobility-triggered routing failures. Due to reliability concerns and the characteristics of clinical environments, we deployed sufficient relay nodes to cover the area in which a user may move. As a result, a mobile node is always one hop away from a relay node. In this case, the impact of mobility is limited to the first hop, i.e.,

the delivery of data from the mobile node to a relay. This enables us to avoid reconstructing end-to-end routes, facilitating a simple and efficient mechanism for reliable data collection from resource-constrained sensors attached to mobile users. DRAP features a mechanism that enables a mobile node to dynamically associate with relay nodes based on channel conditions or in response to mobility.

As part of the CodeBlue project⁵ a publish/subscribe system that uses the TinyAMDR protocol was proposed. TinyAMDR is an adaptation of the AMDR (Adaptive Demand-Driven Multicast Routing) protocol¹¹ to WSNs. TinyAMDR handles mobility by restarting the route discovery process every 15 seconds. Thus, TinyAMDR suffers from the same issues as the routing protocols for ad hoc networks.

A number of projects propose to use body sensor networks for different medical applications such as: assisted living¹², disaster response⁸, and wireless clinical monitoring⁵. These systems emphasize architectural concerns, hardware and medical sensors choice, and specific application requirements. Unfortunately, those projects did not provide systematic study of their networking performance and reliability in the presence of mobile users, which is the focus of this paper.

System Description

Our wireless clinical monitoring system consists of three types of nodes: a base station node, relay nodes, and patient nodes. The base station node is the end-point of the data collection system. The base station node is the only node that requires access to the wired network. The base station stores the received data in a PostgreSQL relational database.

The relay nodes are used to route packets from patient nodes to the base station when there is no direct wireless link between a patient node and the base station. Sufficient stationary relay nodes are deployed to cover the area in which users may move. Relay nodes are connected to power outlets to reduce maintenance cost and enhance reliability. In contrast to existing commercial telemetry systems, we do not require that the relay nodes have access to the wired network, thus removing the high wiring cost. Instead,

the relay nodes form a multi-hop mesh network that covers the area of interest.

Patients carry the patient nodes (Figure 2). A patient node integrates an embedded platform that provides computation, storage, radio, and sensors. The patient nodes are battery operated.



Figure 2. Hardware prototype of a patient node.

Baseline Software Prototype: The goal of the baseline software prototype was to reuse components provided by the TinyOS operating system to develop a clinical monitoring system. The prototype uses the Collection Tree Protocol (CTP)⁹ which is the *de facto* data collection protocol for multi-hop WSNs. CTP is designed to be robust and efficient: it provides close to 90% packet reliability for stationary networks⁹. However, as shown in the next section, CTP performs poorly in the presence of mobility.

Empirical Study

The empirical study is designed to compare the reliability of the baseline prototype in an indoor environment that emulates a clinical environment with mobile patients. The experiments were conducted on a WSN testbed deployed in Jolley and Bryan Halls at Washington University in St. Louis (Figure 1). The testbed consists of 30 TelosB nodes located in Jolley Hall. All sensor data collected from sensor nodes are transmitted over the wireless mesh network. It is important to note that the scale of the experiments matches our target application: monitoring a hospital unit. The formed network is highly redundant, many relays receiving packets from the patient node.

For instrumentation, each node is connected to a central server using a wired USB and Ethernet backbone used to collect experimental results without interfering with ongoing wireless transmissions. In some experiments, an additional 7 TelosB nodes were added to provide coverage in Bryan Hall. These nodes are not connected to the USB backbone and logged their data in flash memory, which is inserted in a database upon completion of each experiment. One of the nodes was selected as the base station while the remaining nodes acted as relays. To better characterize the reliability of the system, in this

section we consider the case of a single mobile user carrying a patient node in a pouch around his neck.

To characterize the network reliability we introduce the following metrics. *End-to-end reliability* is the fraction of data packets received by the base station out of the total number of packets generated by the patient node. *First-hop reliability* is the reliability of the packet transmission over the first-hop, i.e., from the patient node to a relay node. First-hop reliability allows us to directly characterize the impact of user mobility on routing reliability.

Normal activity: In the first experiment, we asked the volunteers to (1) go about their normal routine and (2) record the times when they leave their offices and when they return in order to correlate network reliability and mobility. The experiment lasted for 2.5 hours during which a pulse and oxygenation reading was taken every 30 seconds. We chose this data rate because the pulse and oxygenation typically do not vary significantly within 30 seconds.

Figure 3 shows the vital signs collected during the experiment. The background of the figures is color-coded as follows: a white background color indicates that the volunteer was in his office; the gray background indicates the volunteer outside his office. The end-to-end reliability was 82.39%. A careful analysis indicates a strong correlation between user mobility and packet drops, as consecutive packet drops often occurred following periods of mobility. The first-hop reliability in the two experiments was 85.38%. This indicates that most of the packet drops occur during the first-hop transmission and that once a packet was delivered to the first relay, CTP was able to route the packets through the remaining hops with high reliability.

Mobility Stress Test: We hypothesized that the observed packet drops were due to poor routing table management. The patient node running CTP discovers nodes that are within its communication range. When the subject moves sufficiently to break the link to its current parent in the routing tree, CTP will attempt to transmit to the next best node in its routing table. However, since the subject has moved, many of the entries in the routing table may have been invalidated by the subject's movement and, as a result, CTP will attempt to transmit to nodes outside the patient node's communication range.

To test this hypothesis we asked a subject to move according to the following pattern: stay at position A for 3 minutes, move to position B, stay at position B for 3 minutes, and move back to position A (Figure 1). This mobility pattern was repeated five times. Figure 4 shows the reliability of CTP under mobility.




Figure 3. CTP under normal activity.




Figure 4. CTP under stress mobility




Figure 5. DRAP under normal activity.




Figure 6. DRAP under stress mobility.

The white background indicates the user at position A, the dark the user moving, and grey the user at position B. The pattern of packet drops indicates that our hypothesis is correct: packet drops tend to occur following changes in subject's location.

CTP performed poorly under mobility, having an end-to-end reliability of 71.18% even when a packet was transmitted up to 30 times if it remained unacknowledged. The long sequences of packet drops are particularly disconcerting. The first-hop reliability was 72.88%, indicating that 97.67% of the packet drops occurred in the first hop. In contrast, if the packets from the mobile node are delivered reliably to the first relay node, then CTP's reliability is 97.67%. This suggests that the key research question is to develop a more reliable data collection mechanism from the mobile nodes to appropriate relays, which we call the *dynamic relay association problem* in this paper.

Dynamic Relay Association

We developed the *Dynamic Relay Association Protocol* (DRAP) to achieve reliable data collection from mobile users to their first relay nodes. Based on the insights from the empirical study presented in Empirical Study, DRAP is not designed to be an end-to-end routing protocol for replacing CTP. Instead, it dynamically discovers and selects the right relay for the patient node to send its data after the user moves. DRAP complements routing protocols in that the former focuses on reliable data delivery over the first hop (from the patient node to the first relay), while the latter focuses on delivering data from the first relay to the base station. DRAP can therefore be

easily integrated with existing data collection routing protocols such as CTP.

The key design challenge in DRAP is how to optimize for the common case when there is no mobility and to adapt to user mobility. DRAP accomplishes this by associating a mobile node with the relay node that has the best link quality. DRAP detects when its relay table becomes outdated due to mobility by keeping track of the number of broken links. We consider a link to be broken when five packet transmissions remain unacknowledged. When the number of *consecutive broken links* exceeds a threshold set to three in our experiments, the entire relay table is flushed and new relay is discovered to avoid the previously discussed problem. A detailed description of the DRAP protocol may be found in¹⁰.

Empirical Evaluation: To evaluate the reliability of DRAP we reran the mobility experiments presented in the Empirical Study section. Figure 6 shows the reliability of DRAP under mobility stress. During the one-hour experiment, we did not observe any packet drops. In contrast, CTP dropped 28.82% of the data in a similar experiment (Figure 4). Figure 5 shows the reliability of DRAP under normal activity. During the 2.5-hour experiment, DRAP dropped only 3 packets resulting in an end-to-end reliability of 99.33% at a duty cycle of 2.09%.

Multiple User Reliability Study: To test the performance of the system in the presence of multiple users, we monitored 7 healthy volunteers for an hour. The volunteers followed their normal activities while carrying our patient nodes. The end-to-end reliability of our system is shown in Table 1. The end-to-end reliability for each user was between 96-100%. We

also measured the duty cycle of two of the volunteers, which were 2% and 0.2%, respectively. The difference in the duty cycles may be explained by the difference in the mobility patterns of the two users. The user with more frequent movement had a higher duty cycle because DRAP kept the node active to discover new relays.

User	End-to-End Reliability
1	100.00%
2	98.33%
3	100.00%
4	100.00%
5	100.00%
6	96.15%
7	98.59%

Table 1. End-to-end reliability for multiple users.

An important concern is the scalability of the system as the number of patients is increased. While the experiments are limited to 7 users, the low data rates make it feasible to monitor numerous patients. Typical data rates achieved by 802.15.4 are 25-50 kbps. Assuming a 25kbps data rate and excluding other traffic overhead (e.g., due to routing), the proposed system may monitor up to 4100 patients when the pulse and oxygenation are sampled every 30s. The large margin gives us confidence that we may monitor all patients in a hospital unit.

Conclusion

This paper proposes wireless patient monitoring as an enabling technology for early detection of clinical deterioration in general hospital wards. We present an empirical study on the reliability of real-time data collection from mobile users wearing wireless pulse-oximeters. We observe that the key challenge in such systems is to prevent packet loss from mobile users to their first-hop relays, while the standard data collection routing protocol called CTP can provide reliable data transport over stationary wireless relay networks. Based on this insight we developed the Dynamic Relay Association Protocol (DRAP), an effective mechanism for dynamically discovering right relays for mobile users. The key advantage of DRAP is that it enables reliable data collection from mobile. Empirical evaluation showed DRAP delivered at least 96% of the data from multiple users, while maintaining a radio duty cycle below 2.8%. Our results demonstrate the feasibility of WSNs for clinical monitoring.

This publication was made possible by Grant Number ULI 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 provided by NSF NeTS-NOSS Grant CNS-0627126 and CRI Grant CNS-0708460.

References

1. The Joint Commission. 2008 National Patient Safety Goals. 2008.
2. Kho A, Rotz D, Alrahi K, Cárdenas W, Ramsey K, Liebovitz D, Noskin G, Watts C. Utility of commonly captured data from an EHR to identify hospitalized patients at risk for clinical deterioration. *AMIA Annu Symp Proc.* 2007 Oct 11:404-8.
3. Jones AE, Brown MD, Trzeciak S, et. al. The effect of a quantitative resuscitation strategy on mortality in patients with sepsis: a meta-analysis. *Crit Care Med.* 2008 Oct;36(10):2734-9.
4. Zeitz K, McCutcheon H. Policies that drive the nursing practice of postoperative observations. *Int J Nurs Stud.* 2002 Nov;39(8):831-9.
5. Lorincz K, Malan DJ, Fulford-Jones TRF, Nawoj A, Clavel A, Shnayder V, Mainland G, Welsh M. Sensor Networks for Emergency Response: Challenges and Opportunities. *IEEE Pervasive Computing.* 2004 Oct;3(4):16-23.
6. Hong X, Xu K, Gerla M. Scalable routing protocols for mobile ad hoc networks. *IEEE Network.* 2002 Jul;16(4):11-21.
7. Gao T, Greenspan D, Welsh M, Juang R, Alm A. Vital signs monitoring and patient tracking over a wireless network. *Conf Proc IEEE Eng Med Biol Soc.* 2005;1:102-5.
8. Curtis DW, Pino EJ, Bailey JM, Shih EI, Waterman J, Vinterbo SA, Stair TO, Gutttag JV, Greenes RA, Ohno-Machado L. SMART: an integrated wireless system for monitoring unattended patients. *J Am Med Inform Assoc.* 2008 Jan;15(1):44-53
9. Gnawali O, Fonseca R, Jamieson K, Levis P. CTP: Robust and Efficient Collection through Control and Data Plane Integration. Stanford Univ. Technical Report SING-08-02.
10. Chipara O, Brooks C, Bhattacharya S, Lu C, Chamberlain R, Roman GC, Bailey TC. Reliable Data Collection from Mobile Users for Real-Time Clinical Monitoring. Washington University in St. Louis Technical Report.
11. Jetcheva JG, Johnson DB. Adaptive demand-driven multicast routing in multi-hop wireless ad hoc networks. *Conf. Proc. MobiHoc.*2002;33-44.
12. Wood A, Virone G, Doan T, Cao Q, Selavo L, Wu Y, Fang L, He Z, Lin S, Stankovic J. ALARM-NET: Wireless Sensor Networks for Assisted-Living and Residential Monitoring. Technical Report. Univ. of Virginia CS-2006-11
13. Jon Adams, Bob Heile. Busy as a ZigBee. *IEEE Spectrum.* <http://www.spectrum.ieee.org/computing/networks/busy-as-a-zigbee>