On the Use of Wireless Network Technologies in Healthcare Environments

Nicolas Chevrollier Nada Golmie
National Institute of Standards and Technology
Gaithersburg, Maryland 20899

***Abstract*— In this article, we investigate the suitability of wireless technologies in healthcare/hospital environments. We focus on Wireless Personal Area Network technologies, namely, Bluetooth and the low-rate specifications described in the IEEE 802.15.4 standard. We evaluate the relevance of each technology for supporting medical applications and examine related scalability issues. Moreover, we consider heterogeneous wireless technology environments and quantify the interaction between Bluetooth devices and IEEE 802.15.4 devices when they operate in the same environment.**

I. Introduction

As many hospitals are faced today with increasingly higher wiring cost to plug more devices on the network, there is a practical opportunity to replace wires by wireless technologies. This approach offers significant benefits in terms of reducing deployment costs and providing safer care to patients. However, prevailing over wires by switching to wireless technologies requires a careful analysis of some of the candidate technologies available in order to find out which ones are best suitable for such demanding environments.

As a step in this direction, the IEEE 1073 group is currently developing guidelines for using wireless technologies for medical device communications in various healthcare environments. In fact, from a patient's hospital bedside to a doctor's office, there is a wide range of potential applications and use case scenarios. Medical applications such as real-time waveform delivery, alarm notifications, asset tracking or e-prescription have very strict requirements in terms of accuracy or latency as data lost or delayed have life and death implications but usually have very low data rates. Sensors carrying these applications may be deployed in high density on a patient's body and at the patient's bedside. Other types of applications can also be found in the clinical domain. Queries to hospital databases and Internet access require a fully deployed network infrastructure connecting different departments or hospitals and stress the need for high-speed links to carry bandwidth-hungry applications.

The many constraints imposed by the variety of applications and use case scenarios make the choice of a single fit-all wireless technology difficult if not impossible. Therefore, it is expected that many wireless technologies will have to be used in order to support different application requirements.

In this article, as we focus primarily on low-rate medical applications deployed at a patient's bedside, we consider two potential candidates, namely, the emerging low-rate Wireless Personal Area Network technology as specified in the IEEE 802.15.4 standard [1] and the Bluetooth [2] technology for cable replacement and short range connectivity. We try to answer the following fundamental questions. What are the protocol parameters that are used in mapping medical applications onto wireless technologies? What are the parameter choices that would make this mapping optimal? How scalable are the wireless technologies chosen and how well can they support multiple sensors used on a patient's body?

Most likely, multiple wireless technolgies will be used simultaneously in the same area. As they share the same RF spectrum, the interference level between them is a matter of concern in such unforgiving environments. Thus, after evaluating technologies independently, we investigate whether they can coexist by quantifying the impact of any potential interference.

The remainder of this paper is organized as follows. Section II gives a brief overview of the two selected potential wireless technologies and describes the characteristics and requirements of several medical applications. Section III considers wireless technologies separately and focus on their scalability. In Section IV, we examine a heterogeneous wireless technology environment and provide an evaluation on the impact of interference on performance. The final section offers concluding remarks and future research directions.

II. Wireless Technology Candidates and Medical Applications

In this section, we describe two potential wireless technology candidates for medical applications and give a brief overview of the characteristics and requirements of these applications.

A. Bluetooth

The Bluetooth technology is considered a Wireless Personal Area Network (WPAN), intended for cable replacement and short distance ad hoc connectivity. Bluetooth operates in the ISM frequency band starting at 2.402 GHz and ending at 2.483 GHz in the USA. 79 RF channels of 1 MHz width are defined. The raw rate is defined at 1 Mbit/s and a Time Division Multiplexing technique divides the channel into 625 μs slots. Transmission occurs in packets that occupy an odd number of slots (up to 5). Each packet is transmitted on a different hop frequency with a maximum hop frequency rate of 1600 hops/s.

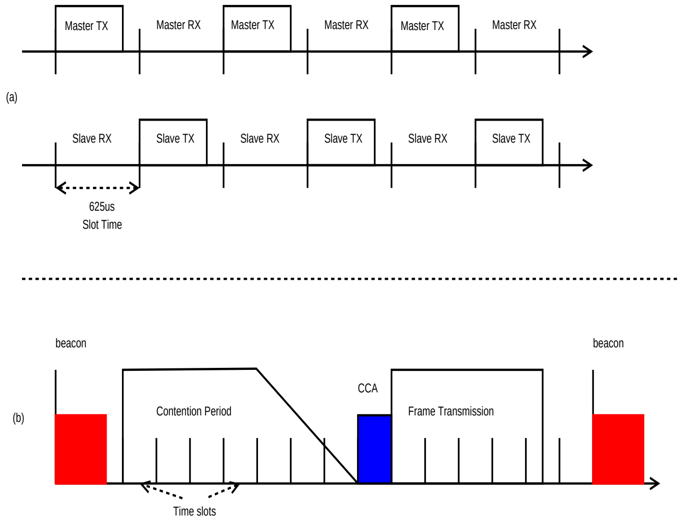

Fig. 1. WPAN structure: (a)Bluetooth (b)Slotted 802.15.4

According to the specifications, up to 8 bluetooth devices can actively participate in a small network, a so-called piconet. Communications inside a piconet occur between a unique predominant device, the master, and subordinate devices, so-called slaves. Upon connection establishment, a slave synchronizes its time and frequency hopping to the master's and waits to be polled by the master to transmit. In this manner, a slave's packet always follows a master's packet. Figure 1 describes the transmission of one-slot type packet between a master and one slave.

B. IEEE 802.15.4

Another candidate for carrying medical applications is the low rate IEEE Std. 802.15.4-2003. Designed for low cost products, it supports very limited battery consumption, a short range operation (10m) and low rate applications. Two different variations of the technology can be found: the slotted and the unslotted channel structure.

- The slotted channel structure described in Figure 1 uses synchronization between devices enabled by beacon exchanges and a slotted CSMA/CA mechanism as described in [4]. When a device wishes to send data frames, it waits for a random number of slots. Then, the medium idleness is evaluated during a CCA (Clear Channel Assessment) period of time. If the medium is still idle at the end of this period, the packet can be sent at the beginning of the next time slot, otherwise the procedure is restarted from the beginning.
- In the unslotted channel structure, if a device wishes to send data frames, it waits for a random period of time. Then if the medium is still idle after a CCA period of time, the frame transmission can start.

The physical layer describes three different frequency bands:

- 1 channel in the (868 to 868.6)MHz band providing 20 kbit/s
- 10 channels in the (902 to 928)MHz band providing 40 kbit/s each
- 16 channels in the (2400 to 2483.5)MHz band providing 250 kbit/s each

We focus our effort on the unslotted version in the 2450 MHz band as it provides the highest data rate combined with the least overhead (i.e., no beacon frames).

C. Application Requirements

In this section, we describe the nature of some medical applications and their requirements that have life or death implications when data is lost, corrupted, or delayed. This is unlike most other environments where these types of requirements are mainly financial.

As part of the framework evaluation, the IEEE 1073 group has defined a number of potential medical applications and usage cases. Each medical application is defined in terms of a data rate (raw data needed to be transported), end-to-end latency (potential packetization and transmission delays), and expected coverage area (radio distance between two communicating devices).

An example of a medical application is the electrocardiogram (ECG) monitoring. It uses a star topology where multiple sensors communicate with a unique collector. An ECG is an electrical recording of the heart used in the investigation of heart disease. It can identify abnormalities in the heart's electrical conduction system. The data stream resulting from the digitized analog signal is sent to a control monitor, available on either a nurse's personal digital assistant (PDA) or a nurse's personal computer (PC). As part of an ECG system, a Personal Worn Device (PWD) defined by the IEEE 1073 group (i.e., a wireless electrode) generates 4 kbit/s of data and requires that the addition of the latency introduced by the packetization of the samples and the transmission delay remain below 500 ms.

The goal of our evaluation is to determine how well Bluetooth and IEEE 802.15.4 support the IEEE 1073 medical application described above and identify any scalability issues.

III. Scalability Issues with Medical Sensor Deployment

In this section we focus on the scalability issues pertaining to the use of the Bluetooth and the IEEE 802.15.4 wireless technologies for medical sensors. We compare the performance obtained with each technology to support an ECG application.

A. Topology and Simulation Parameter Setting

We focus our attention on the ECG system which requires the deployment of multiple electrodes on a patient's body, each of them carrying a low data rate application. We use the topology depicted in Figure 2. The distance between the communicating devices remains within the constraints of a room. Since up to 16 leads can be used on a patient's body, this may represent a scalability issue for the technology considered.

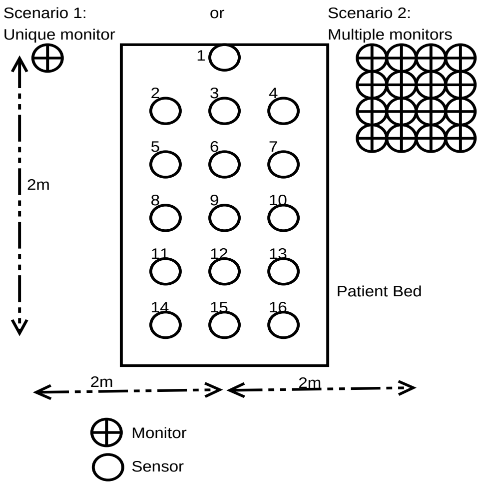

Fig. 2. Patient bed topology

In addition to the configuration requirements, the medical application data stream needs to be formatted into packets and mapped onto the baseband framing available. Among the many choices available, we try to minimize the overhead while using the full payload of each packet.

In the case of the IEEE 802.15.4 technology, 944 bits per packet are available at the application layer. Thus, to achieve at least the minimum rate required of 4kbit/s, a packet has to be generated every 0.236s.

In the case of the Bluetooth technology, we can choose either an Asynchronous Connection-Less or a Synchronous Connection-Oriented (SCO) link. The latter is a symmetric point-to-point connection between a master and a specific slave where a master sends an SCO packet to the slave at regular time intervals. The time interval can either be 2, 4 or 6 time slots for HV1, HV2, or HV3, respectively. SCO links are primilary designed to support voice traffic, thus we use an ACL link, intended for data communications. An ACL link is an asymmetric point-to-point link between a master and active slaves in the piconet using retransmissions to ensure data integrity. Among all available packet formats on an ACL Link (i.e., DM1, DM3, DM5, DH1, DH3, DH5), DM5 and DH5 are not pertinent as the amount of data carried by a full DM5 or DH5 packet (i.e., 1760 bits and 2680 bits respectively) implies a significant packetization delay. DH3 and DM3 are the next available framings as we want to minimize overhead. We select DM3 as its payload size of 936 bits is very similar to the 944 bits available in an IEEE 802.15.4 packet, which makes a comparison between the use of these two technologies more practical.

Given the 625 μs time slot and the 936 bits payload size of a DM3 packet, a device needs to be polled every 374 slots to achieve at least the minimum rate required of 4 Kbit/s for the PWD application. The interarrival between two packet generations is then 0.23375s (374*625μs) which represents the packetization delay.

We develop models for both technologies using the com-

TABLE I
SIMULATION PARAMETERS

	802.15.4 Sensor	Bluetooth Sensor
Transmitted power (mW)	1	1
Packet header (bit)	72	174
Payload size (bit)	944	936
Packet interarrival time (s)	0.236	0.23375

mercial network simulation package OPNET[1]. Our simulation environment is based on detailed MAC, physical layer (PHY) and channel models. The parameters used in the simulations to model an electrode are summarized in Table I.

B. Simulation Results

We use multiple parameters to evaluate performance of a given scenario. The performance metrics include:

- End-to-end delay as presented in Figure 3.
- Packet loss at the MAC sublayer of the receiver node (i.e., the monitor) as shown in Figure 4.
- Efficiency, representing the number of successful data packets received at the receiver's application layer divided by the number of data packets generated by all the transmitter's application layers related to this receiver as plotted in Figure 5.

In order to comply with medical requirements, the end-to-end delay combined with the latency introduced by the packetization has to be below 500 ms and the efficiency has to be equal to 1. Every application layer packet generated has to be transmitted and received.

The topology depicted in Figure 2 includes two scenarios. In scenario 1, a single access point is used for the central monitor in order to collect data from all devices placed on the patient's body. When using a single monitor, serious limitations exist in terms of scalability, especially for Bluetooth. In fact, due to the protocol specifications, only 7 slaves can be part of a single piconet, thus allowing only 7 sensors or electrodes to be deployed on the patient's body. In this case, as sensors have to be polled every 374 slots, there is enough bandwidth to acccomodate all 7 slaves. The application delay increases gradually to reach 0.00975s when 7 slaves transmit data to the central monitor. This is only due to the round robin mechanism used to poll alternatively each slave.

To overcome the protocol's limitations in Bluetooth, another option is to use a different piconet per sensor/central monitor pair. This is referred to as scenario 2, in Figure 2, where each lead uses a different piconet to send data to its own central monitor. In this case, the interference resulting from multiple Bluetooth piconets operating in close proximity may lead to a higher packet loss at the central monitors as seen in Figure 4. In Figure 5, we see that up to 13 sensors,

[1]Certain equipment, instruments, or materials are identified in this paper in order to specify the experimental procedure adequately. Such identification is not intended to imply recommendation or endorsement by the National Institute of Standards and Technology, nor is it intended to imply that the materials or equipment identified are necessarily the best available for the purpose

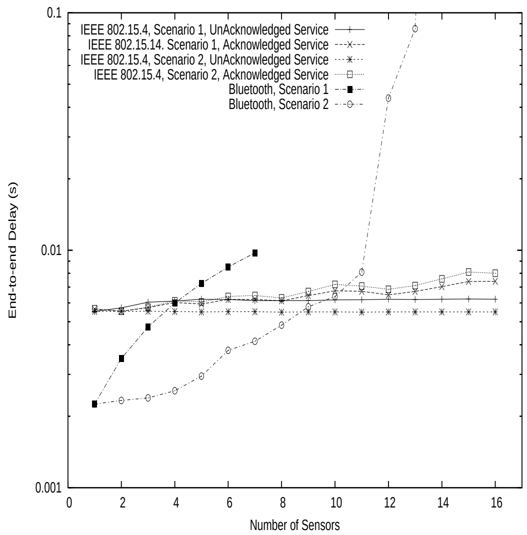

Fig. 3. Delay as a function of number of sensors

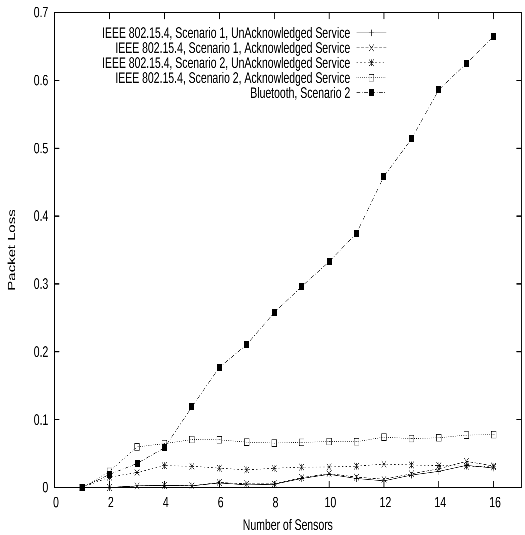

Fig. 4. Packet Loss as a function of number of sensors

the efficiency equals 1 with an end-to-end delay (Figure 3) that combined with the packetization latency (i.e., packet interarrival) does not exceed the 500ms limit. Each data point in the curve labeled "Bluetooth, Scenario 2" is an average over the number of devices considered. Beyond 13 piconets, the packet loss in Figure 4 is so high that the efficiency decreases dramatically and the additional of the end-to-end delay and the packetization latency exceeds the required limit. Thus, for this specific scenario, up to 13 electrodes can be supported by dedicating a piconet to a pair of electrode-monitor. Another design we can try is to use multiple piconets, each of them carrying data from multiple electrodes. For example, at least 3 piconets can be used to support 16 electrodes. In this case, 10 electrodes can be split on two piconets (5 on each), while the third piconet will have 6 electrodes. When all 16 devices are running at the same time, an efficiency of 1 can be achieved. Although feasible, the main disavantage of this approach lies in the added configuration and deployment complexity.

We repeat the same two scenarios using IEEE 802.15.4. In scenario 1, multiple IEEE 802.15.4 sensors communicate with a single access point or central monitor. For the unacknowledged mode in Figure 5 the efficiency starts to drop when 3 or more devices belong the same network. There are two explanations for this phenomenom. In fact, even if devices start their transmission randomly at the beginning of each simulation, at some point they might end up trying to access the medium at the same time which leads to packets colliding at the receiver. In this unacknowledged mode, packets that are lost are not retransmitted. Figure 4 shows a significant packet loss when using more than 8 devices. The second explantation is that even if a transmitter senses the medium to be busy, it has only 4 attempts to access the medium. Between each attempt, a random backoff has to be performed. As we use fairly long packets compared to the possible backoff window, these attempts can be unsuccessful resulting in dropping a

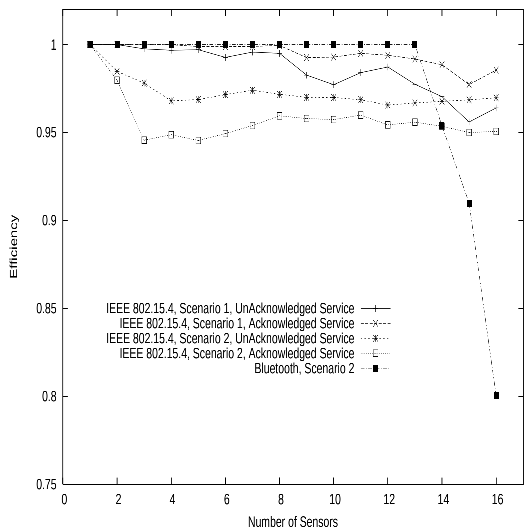

Fig. 5. Efficiency

packet on the transmitter side. Even if the drop rate is small (about 1% when 3 or more sensors are used), it is a major disadvantage when using medical applications with strict loss requirements.

Using the acknowledged service solves the issue of packet collisions by retransmitting lost packets but does not provide any solution against packets dropped at the transmitter. In fact, packets can also be dropped at the transmitter side when 3 transmissions have failed to receive a proper acknowlegment. We notice that the drop rate is about 2% when more than 3 sensors are used. The percentage of packet loss is slightly higher than in the case of the unacknowledged service since more packets are exchanged due to (multiple) retransmissions.

For scenario 2, we set each a pair of electrode-monitor on a different channel using up all 16 channels available. First, we run the unacknowledged service. We realize that, for our particular topology, there is a significant packet loss as shown in Figure 4 due to co-channel interference making this service not suitable for medical applications. In addition, using the acknowledged service does not improve performance. As packets are retransmitted, more packets are exchanged resulting in a higher packet loss. Efficiency does not improve either. In fact, a transmission on a specific channel is unaware of transmissions on adjacent channels. Therefore, when two packets are about to be sent on adjacent channels, there is no mechanism to avoid transmission overlap and packet collisions. In this case, retransmissions will occur almost simultaneously and considering the significant packet size, packets will most likely collide again. After 3 tries, packets will eventually be dropped at the transmitter's side.

Considering our particular topology and our specific mappings of medical applications onto wireless technologies, it appears that scalability is not guaranteed. Both Bluetooth and the IEEE 802.15.4 technologies have severe limitations for deploying multiple sensors on a patient's body. On one hand, protocol specifications and interference strongly limit the deployment of Bluetooth sensors and on the other hand, limited bandwith and MAC protocol design limits the use of 802.15.4 equipments. Moreover, by using these two technologies, the 2450 MHz frequency band is occupied and most likely it will not be interference-free.

IV. Coexistence of Bluetooth and IEEE 802.15.4 in the same environment

While the first part of this article focuses on scalability issues for the Bluetooth and the IEEE 802.15.4 technologies and their deployment in a medical environment, we turn our attention next to investigating how well they can coexist in the same environment. In fact, the deployment of wireless technologies has already started in many hospitals. From internet access to file transfer, WLAN is used heavily in healthcare environments. Nurses carrying PDAs equiped with Bluetooth connectivity exchange patients' information on a regular basis. As time progresses, we can envision using IEEE 802.15.4 sensors on a patient's body in order to collect critical information and sending it to a central monitor located at the patient's bedside.

Interference between Bluetooth and WLAN devices and its impact on performance has been well documented in the literature [6][7] and coexistence solutions have been proposed [8][9]. In addition to these evaluations and coexistence schemes, Golmie *et al.* [5] started to examine the interaction between WLAN and 802.15.4 devices. They noticed that a WLAN device can significantly impact IEEE 802.15.4 devices. In some experiments presented in [5], communication between IEEE 802.15.4 devices was simply not possible. Nevertheless, both IEEE 802.15.4 and WLAN use spread spectrum techniques and one could argue that coexistence between these technologies is only a matter of choosing adequately non-overlapping channels. Consequently, in this article, we focus our effort on evaluating the interactions between Bluetooth and IEEE 802.15.4 devices when they are present in the same environment.

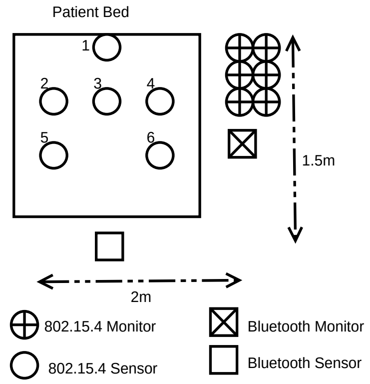

Fig. 6. Heterogeneous Wireless Environment

A. Topology and Simulation Setting

We extend scenario 2 described in Figure 2 by mixing different technologies in the same environment. Six IEEE 802.15.4 sensors (i.e., sensor1, sensor2, sensor3, sensor4, sensor5, sensor6) are spread over a patient's body according to the topology shown in Figure 6 and transmit medical information following the description in Table I. To avoid any interference between the IEEE 802.15.4 devices, non-overlapping channels are used and carriers are set 15MHz apart, thus using 6 channels. We use the acknowledged service which can potentialy overcome packet loss. Meanwhile, in the same room, a nurse carrying her PDA sends information via a Bluetooth connection to a Bluetooth access point located close to the IEEE 802.15.4 monitors (All monitors are located close to the patient's bedside less than 0.5 meters away). We use the definition of the PDA application from the IEEE 1073 group which requires 60kbit/s to be sent. Using DH3 framing, a packet is sent every 0.02375s. As we add, IEEE 802.15.4 sensors one by one, we examine the interactions between the two technologies.

B. Simulation Results

First, we look at the impact of a Bluetooth transmission on the IEEE 802.15.4 communications. Figure 7 shows the average packet loss for different IEEE 802.15.4 sensors labeled ”1” to ”6”. Differences between sensors are significant due to their location and position with respect to the Bluetooth transmitter. Thus, each is impacted differently by the Bluetooth transmission. Nevertheless, a factor remains constant; they are all severely impacted as the packet loss ranges between 26% for sensor6 to 62% for sensor5. During the transmission of a single IEEE 802.15.4 packet, the Bluetooth device would have hopped on 10 different frequencies, causing errors in the IEEE 802.15.4 packet being received that then must be dropped. For most of the sensors, using the acknowledged service does not overcome packet loss as retransmissions suffer the same fate as the initial transmission. Thus, efficiency drops below 1 to reach at most 0.83 for sensor6.

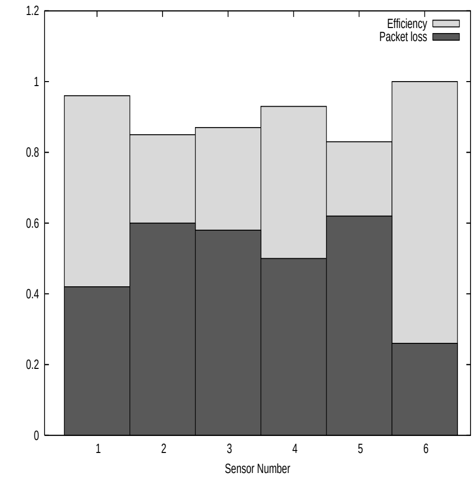

Fig. 7. Packet Loss and Efficiency per IEEE 802.15.4 Sensor

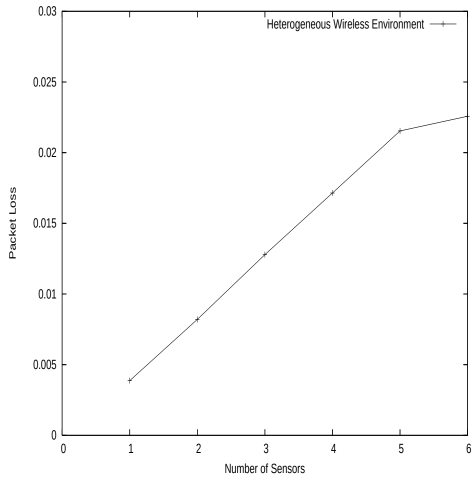

Fig. 8. Bluetooth Packet Loss as a function of the number of IEEE 802.15.4 sensors

On the other hand, Bluetooth devices are relatively less impacted. Figure 8 shows the packet loss recorded at the MAC layer of the monitor. As more IEEE 802.15.4 sensors are added, the packet loss increases steadily to reach about 2.5%. There are several explanations for this behavior. In this case, the choice of the topology plays an important role in the simulation results. The IEEE 802.15.4 sensors are relatively far from the Bluetooth monitor in order to cause significant interference. On the other hand, the IEEE 802.15.4 monitors that are closer to the Bluetooth monitor and could cause more interference, mostly receive data sent by the sensors. In the acknowledged service, the IEEE 802.15.4 monitors return ACK packets that are relatively short and do not cause significant packet loss on the Bluetooth monitor as seen in Figure 8.

V. Conclusion and Future Work

In this article, we investigate the use of two wireless technologies, namely, Bluetooth and IEEE 802.15.4, in healthcare environments. Our findings are summarized as follows.

First, the scalability of these technologies is not a given feature. Both Bluetooth and IEEE 802.15.4 have major constraints in terms of supporting topologies consisting of multiple medical sensors. The protocol specifications (i.e., the limitation of 7 slaves in a Bluetooth piconet) and the interference between multiple piconets significantly limit the use of Bluetooth for medical sensors. Meanwhile, limited bandwidth and MAC protocol design restrict the use of the IEEE 802.15.4 technology. Very specific topologies need to be carefully designed in order to support a high sensor density area using either one of these two technologies.

Moreover, both technologies use the same RF spectrum and using them simultaneously leads to severe interference and performance degradation. In our experiments, 802.15.4 devices are strongly impacted by a nearby Bluetooth communication. Our results show that these technologies are unable to meet very strict application requirements under certain assumptions chosen in this paper and thus, their usage in a healthcare environment may require careful configuration design and even protocol enhancements. Future simulations using different assumptions (i.e., packet size) will help us to quantify a trade-off between packet loss, latency and overhead.

Prevailing over wires in healthcare environments by using wireless technologies implies searching for mechanisms to overcome interference between different technologies. For the Bluetooth technology, we will adapt existing mechanisms such as Adapative Frequency Hopping (AFH) [7], originally developed to mitigate interference between WLAN and Bluetooth, in order to enable coexistence between IEEE 802.15.4 and Bluetooth devices. In the meantime, our plan is to run complex scenarios with multiple wireless technology devices operating simultaneously in the same area to explore the impact of interference among them.

References

[1] IEEE Std. 802.15.4-2003, ”Standard for Telecommunications and Information Exchange Between System - Local Area Metropolitan Area Networks - Specific requirements - Wireless Medium Access Control (MAC) and Physical Layer (PHY) Specifications for Low Rate Wireless Personal Area Networks (WPAN),” 2003.

[2] Bluetooth Special Interest Group, ”Specifications of the Bluetooth System, vol. 1, v.1.0B ’Core’ and vol. 2 v1.0B ’Profiles’,” Decembre 1999.

[3] IEEE Std. 802-11, ”IEEE Standard for Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specification,” June 1997.

[4] G. Lu, B. Krishnamachari and C. Raghavendra, ”Performance Evaluation of the IEEE 802.15.4 MAC for Low-Rate Low-Power Wireless Networks,” in *IEEE IPCCC EWCN Workshop*, April 2004.

[5] N. Golmie, D. Cypher and O. Rebala, ”Performance Analysis of low rate wireless technologies for medical applications,” in *to appear in Elsevier Computer and Communications*, 2005.

[6] N. Golmie, N. Chevrollier, and O. Rebala, ”Bluetooth and WLAN Coexistence: Challenges and Solutions,” in *IEEE Wireless Communications Magazine*, Vol. 10, No. 6, December 2003.

[7] N. Golmie, R. E. Van Dyck, A. Soltanian, A. Tonnerre, and O. Rebala, ”Interference Evaluation of Bluetooth and IEEE 802.11b Systems,” in *ACM Wireless Networks*, Vol. 9, pp. 202-211, 2003.

[8] N. Golmie, ”Bluetooth Dynamic Scheduling and Interference Mitigation,” in *ACM Mobile Networks*, MONET Vol. 9, No. 1, February 2004.

[9] Carla F. Chiasserini, and Ramesh R. Rao, ”Coexistence mechanisms for interference mitigation between IEEE 802.11 WLANs and Bluetooth”, in *Proceedings of INFOCOM 2002*, pp. 590-598.