Design and Implementation of Multi-User Wireless Body Sensor Networks

José A. Afonso, Pedro Macedo, Helder D. Silva, José H. Correia, Luis A. Rocha
Department of Industrial Electronics
University of Minho
Guimarães, Portugal
{jose.afonso, pmacedo, dsilva, higino.correia, lrocha}@dei.uminho.pt

***Abstract*—This paper describes the development of two low power wireless body sensor network (BSN) prototypes that allow the real-time monitoring of wearable sensors data from multiple users with a single base station. The prototypes present differences in the architecture, wireless network hardware and implemented protocols. These systems integrate inertial sensors that allow the monitoring of users' posture and goniometric development, as well as heart rate and respiratory rate sensors. The design of body sensor networks is challenging due to the requirements of quality of service provisioning and low energy consumption, combined in many cases with high traffic loads. The solutions presented in this paper address these challenges, with particular focus on the implemented medium access control protocols, which overcome some drawbacks presented by standard protocols when applied to BSN scenarios. The performance of the developed systems is presented, compared and discussed using theoretical, experimental and simulation methods. Results show that the proposed mechanisms provide good bandwidth efficiency and decrease the delivery error rate without significant increase in the energy consumption.**

Keywords-body sensor networks, medium access control, wireless sensor networks, quality of service

I. INTRODUCTION

A body sensor network (BSN) is characterized by the presence of sensors nodes attached to a user's body in order to acquire physiological signals and a central monitoring unit, which receives the data from the sensors, either directly or through intermediate nodes. These networks may allow the monitoring of several kinds of signals, such as electrocardiogram (ECG), oximetry, temperature, heart rate, blood pressure and respiratory rate, as well as the motion of segments of the body. Conventional wired BSN monitors have been used in places like hospitals over the last decades; however, these systems impose significant restrictions on the users' mobility.

Recent advances in microelectronics, wireless communications and sensors technology, with the corresponding reduction in size and cost, foresee a rapid expansion on the development of wireless body sensor networks over the next years, enabling mobile and continuous monitoring of users even during their daily life activities, with applications in the areas of healthcare, sports and entertainment.

In the healthcare area, although extensive physiological data monitoring is currently available in many hospitals, the monitoring and diagnostic procedure is normally limited in time, restraining the capability of these systems to properly capture the patients' physiological states, since transient abnormalities cannot always be detected in this scenario. For example, many cardiac diseases are associated with relevant episodic abnormalities such as arrhythmias, episodes of myocardial ischemia or transient surges in blood pressure [2]. Since the timing of these abnormalities is impossible to predict, much time and effort is wasted when trying to capture an episode with controlled monitoring using wired systems, making many vital and even life threatening disorders go undetected due to their random behavior. Wireless body sensor networks facilitate the continuous wireless monitoring of users, having the potential to replace life saving but expensive therapies. It is expected that these systems will be used initially to monitor at-risk patients, with their use being gradually extended to benefit a larger proportion of the population.

Despite of the benefits that can result from the adoption of wireless BSNs, there are some challenges that still limit their widespread application. Normally these systems need to satisfy demanding requirements in terms of quality of service (QoS), such as reliable data delivery, sustainable throughput and bounded delay. At the same time, since sensor nodes are normally battery powered, it is necessary to minimize their energy consumption to increase their lifetime. A challenge arises from the fact that some kinds of sensors must be sampled quite often, generating a great amount of data and, consequently, requiring the network to operate under high loads. For example, a motion capture system based in inertial sensors placed on the body needs to sample multiple sensors several times per second to provide satisfactory results (30 frames per second in many applications). Another example is ECG monitoring, which can require sampling rates of up to 250 Hz per lead with a resolution of 12 bits [3]. If contention-based MAC (Medium Access Control) protocols typically proposed for wireless sensor networks are used, the high traffic load generated by these data-intensive applications will lead to frequent collisions between sensor nodes.

This paper describes the design and development of two multi-user wireless BSN prototypes that enable the simultaneous monitoring of the posture and movements of several users in real-time with a single central monitoring

unit (base station). An application of this systems is the monitoring of teams of athletes in a gymnasium for sports such as basketball, volleyball, handball and gymnastics, with the goal of providing detailed information in order to enhance the performance of both the athletes individually and the team as a whole. Another application is in the medical field, namely in physiotherapy sessions, where such a system can benefit both the patient, by increasing his levels of confidence, and the therapist, by providing detailed information about the patient evolution. The first system prototype also implements sensors to measure the heart rate and respiratory rate.

This paper is an extended version of [1]. Compared to the original paper, it includes a section about wireless network standards and extends the description of the first developed prototype. In addition, this paper also includes the description of a new prototype, comprising an overview of its design characteristics, the presentation of its enhanced MAC protocol and the presentation and discussion of simulation results comparing the performance of the protocols used by both prototypes.

Next section presents an overview of current wireless standards and discusses their suitability to BSN applications. Section III presents the architecture of the first prototype, the implemented MAC protocol (LPRT) and the developed PC application, as well as a theoretical analysis and experimental results. Section IV presents the architecture of the second prototype, describes its MAC protocol (iLPRT) and presents simulation results. Finally, Section V presents the conclusions.

II. Wireless Network Standards

Most wireless networks can be classified in the following five categories: wireless local networks (WLAN), wireless personal networks (WPAN), satellite networks, mobile cellular networks and broadband wireless access (BWA) networks. Some characteristics of former two differ substantially from the characteristics of the latter three. WLAN and WPAN are relatively short range technologies that normally work in unlicensed frequency bands and where all the network equipment normally belongs to the user. On the other hand, the latter networks tend to work in licensed bands, with the infrastructure belonging to the network operator, which charges the user for the provided services. The characteristics of the former are more convenient for the applications considered in this paper, so those networks will be discussed further in the remaining of this section.

The most common WLANs available in the market are based on the IEEE 802.11 standards [4]. WLANs present higher range and data rates than most WPANs. However, they normally present also higher energy consumption and large protocol overhead with short data packets, which make them less adequate to be used with battery-operated body sensor network nodes. For example, the MAC header and trailer of 802.11 data frames occupy a total of 34 bytes, which represents a significant overhead given that the payload generated by sensor applications is normally composed of few bytes.

Bluetooth [5] is a WPAN technology that works in the 2.4 GHz band using frequency hopping spread spectrum (FHSS). The network, called piconet, is organized in a star topology composed by one central node (master) and a maximum of seven active end nodes (slaves). A piconet has a range from 10 to 100 m, depending on the class of the device. Piconets can be interconnected to form a scatternet. Bluetooth uses a polling based MAC protocol that provides support for both real-time and asynchronous traffic. However, the restriction of only seven slaves per piconet imposes a severe limitation on the number of wireless nodes of the BSN system. While the scatternet functionality provides a way to increase the number of the nodes of the network, its multi-hop multi-channel and single radio architecture tends to increase the complexity and decrease the performance of the network, making it harder to provide QoS guarantees to the application. Another obstacle for the implementation of large scale BSNs based on Bluetooth is that few commercially available devices are able to support the seven slaves per piconet specified by the standard, and even less, if any, are able to properly support scatternets.

The standardization of WPAN networks at the IEEE is under the scope of the 802.15 group. One of these standards, the IEEE 802.15.4 [6] defines the physical and MAC layer of ZigBee [7], which targets wireless sensor network (WSN) applications through the provision of low power and low bit rate WPANs. At the physical layer, the IEEE 802.15.4 relies on direct sequence spread spectrum (DSSS) to enhance the robustness against interference, providing gross data rates of 20/40 kbps, at the 868/915 band, and 250 kbps, at the 2.4 GHz band. The basic IEEE 802.15.4 MAC protocol is a contention-based CSMA/CA (Carrier Sense Multiple Access with Collision Avoidance) mechanism, which is not the most suitable option to provide the QoS required by BSN applications, especially under high loads, due to the increased probability of collisions. The performance of the CSMA/CA mechanism is also negatively affected by collisions due to the hidden node problem.

The IEEE 802.15.4 standard also provides a guaranteed time slot (GTS) scheme in order to support devices requiring dedicated bandwidth and low latency transmission. However, only seven GTS allocations per superframe are supported with this scheme, imposing a limitation on the number of wireless nodes of the BSN system. Besides that, the granularity of the transmission periods that can be allocated to the nodes is coarse, since each superframe is composed by just 16 slots, which can result in wasted bandwidth, reducing the bandwidth efficiency.

Given the drawback of the current wireless networks standards for the BSN applications considered in this paper, we decided to design and implement our own MAC protocols. These implementations are based on low power COTS (Commercial Off-The-Shelf) wireless network components that are compliant with the IEEE 802.15.4

standard at the physical layer, in order to benefit from the economies of scale (low cost) and component integration (small size).

III. First System Prototype

An emerging new area of research that combines the strengths and capabilities of electronics, sensors and textiles is the field of electronic textiles (e-textiles) [8]. If these e-textiles are further combined with wireless communication capabilities, wireless smart textiles systems with local monitoring and computation skills and remote data storage become reality [9].

A. Prototype Overview

The architecture of the first multi-user wireless BSN prototype is composed by a personal computer (PC), a base station and several smart shirts (wireless stations), as depicted in Figure 1. Each shirt aggregates data from a heart rate sensor, a respiratory sensor and several posture sensors, which allow the measurement of the user's posture and goniometric development, and transmits this information continuously to the base station using its wireless station. Microcontroller, power supply, RF transceiver and antenna are also attached to the shirt. The base station is connected to a PC, where the data collected from the shirts can be stored and visualized in real-time.

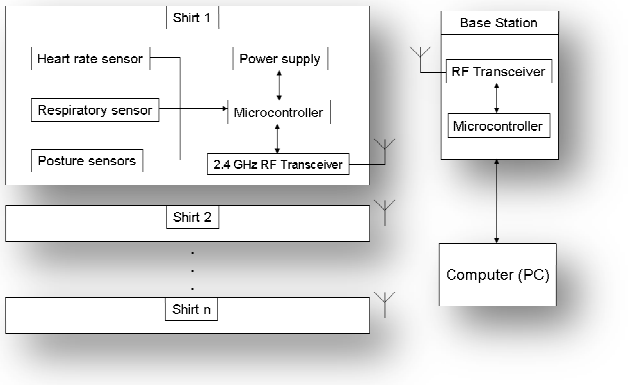

Figure 1. Architecture of the first prototype.

An ear clip infrared sensor is used to measure the heart rate, based on the skin reflectivity to infrared light. Since the reflectivity changes with blood pressure, which in turn changes with the pumping action of the heart, the variations in reflectivity are used to measure the heart rate.

The respiratory frequency is measured with inductance plethysmography, which measures cross-sectional area variation and is commonly used to carry out respiration measurements. The respiratory sensor is made of a copper yarn crocheted in a polyester belt, providing an adjustable stretch due to the addition of rubber yarns. If the belt is sewed on the suit around the abdomen or thorax, circumference and length changes of the belt caused by breathing result in an inductance variation, which can be measured and processed.

The posture monitoring system is composed by one module for each segment of the body requiring monitoring. The developed shirt has sensor modules in the spine, hips and arms. The same modules could also be used to measure the angles of other segments of the body, such as the legs, if desired. Each module contains a 3-axis accelerometer and a 3-axis magnetometer that are used to obtain the pitch, roll and yaw angles. Both the gravitational force and the Earth's magnetic field are used to detect the angles of the segments. The former is used to detect inclination while the later is used to measure the rotation of the body about the axis perpendicular to the gravity field. Figure 2 shows the posture sensor module that was developed. Each module is able to provide angles resolution of around 1 degree. The posture monitoring system is described in patent WO 2008/018810 A2 [10].

Both the base station and the shirt stations are implemented using Crossbow MICAz motes [11] (Figure 3). The MICAz integrates a microcontroller (Atmel ATMega128) and a RF transceiver (Chipcon CC2420) that is compliant with IEEE 802.15.4 at the physical layer.

Each shirt can be seen as a wired BSN where the multiple sensor modules are linked to the MICAz mote through the use of an I^2C serial bus. The shirt electronics is powered by two AA batteries.

The base station is attached to a MIB520 interfacing and programming board, which allows the MICAz to communicate with and receive power from the PC through a USB cable.

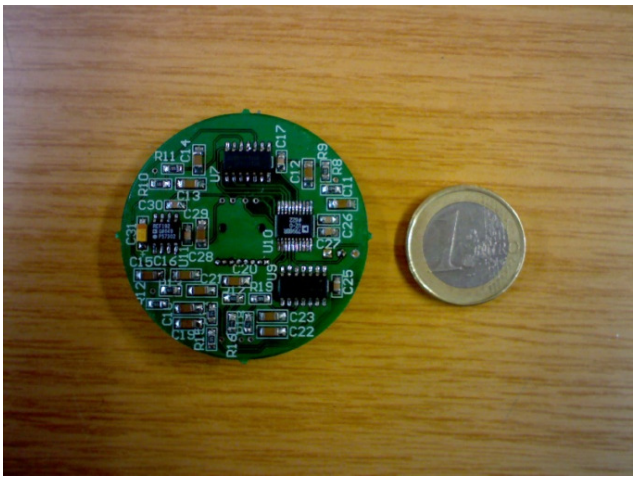

Figure 2. Posture sensor module.

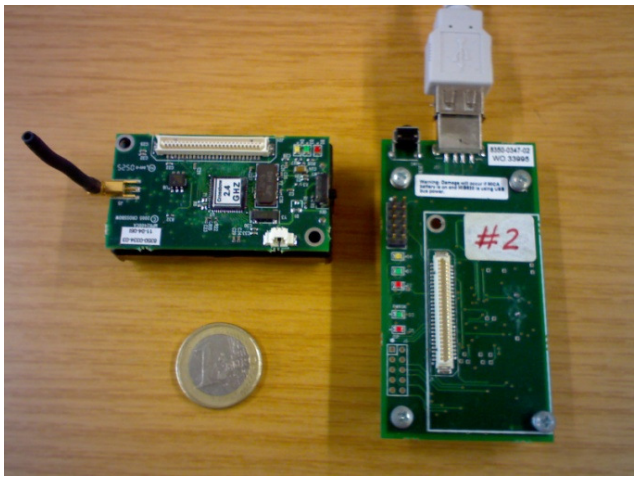

Figure 3. MICAz (left) and MIB520 (right).

B. *The LPRT MAC protocol*

The LPRT (Low Power Real Time) protocol was implemented in the MICAz motes using the TinyOS operating system and the nesC programming language [12]. LPRT is a hybrid schedule-based dynamic TDMA protocol and contention-based CSMA/CA protocol that is based in the superframe structure presented in Figure 4. Each superframe is divided in a fixed number of mini-slots (1024, in this implementation), and starts with the transmission, by the base station, of the respective beacon frame (B), which is followed by the Contention Period (CP). During the CP any station can transmit using the rules of a CSMA/CA protocol. The CP allows the stations to associate with the base station and to request mini-slots for transmission during the Contention Free Period (CFP). It can also be used to convey non-real-time asynchronous traffic. Stations must not initiate a CSMA/CA transaction if it cannot be completed before the beginning of the CFP.

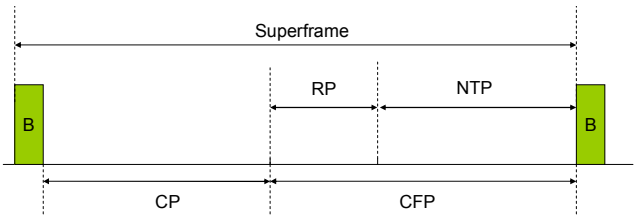

Figure 4. Superframe format of the LPRT protocol.

The Contention Free Period is placed after the CP. Transmissions during the CFP are scheduled by the base station using resource grant (RG) information announced previously in the beacon frame of the current superframe. Since transmissions during the CFP are regulated by the base station, they are not affected by the hidden node problem, unlike transmissions of protocols that rely on carrier sensing.

The CFP is composed by an optional retransmission period (RP) and a normal transmission period (NTP). The RP gives a new opportunity for transmission of data frame when it was not delivered correctly in the previous superframe. The RP is placed after the CP so that a data message is retransmitted away from an eventual burst error condition responsible for the transmission failure occurred during the previous NTP. The RP is placed before the NTP to allow the data message to reach its destination before the next one is transmitted, in order to limit the delay and preserve the sequence of data. This division is not mandatory, however, as retransmissions can be mixed with regular transmissions during the entire CFP period. The retransmission procedure is intended to increase the reliability of the protocol, which is fundamental for applications with low loss tolerance.

MICAz motes use fixed 16 bit long station addresses. Instead of using these addresses in the resource grant (RG) fields of the beacon to identify the stations that are allowed to transmit during the CFP, the LPRT protocol uses a smaller association ID (AID) that is dynamically assigned to the stations when they associate with the base station. This approach reduces the length required for the beacon, making it less sensible to channel errors. The current implementation uses 5 bit long AID addresses, allowing a maximum of 32 associated stations, which is sufficient for the envisioned applications.

Figure 5 shows the structure of the payload of the beacon frame. The superframe duration field gives the duration of the current superframe in multiples of a minimum superframe duration time. It is followed by a list of resource grant (RG) fields, whose quantity is expressed by the RG list length (RLL) field. Each RG is composed by a transmission direction (TD) bit, the association ID (AID) field and an initial transmission slot (ITS) field. The RG allows the scheduling of data transmissions either to or from the station identified by the AID, depending of the value of the TD bit: “0” for downlink and “1” for uplink direction. More than one RG can be granted to a station in the same superframe. The total transmission period granted by a given RG goes from the beginning of respective ITS until before the beginning of the ITS of the next RG on the list. The uplink data can include piggybacked information for alteration of resource allocation parameters, if desired.

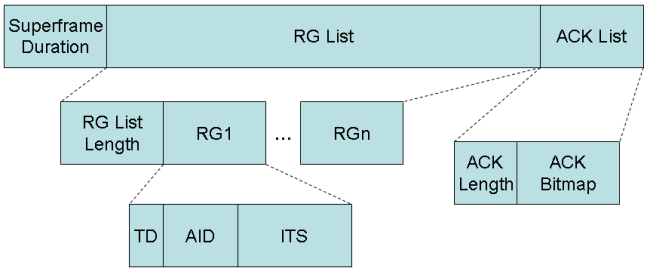

Figure 5. Structure of the beacon frame payload.

For downlink transmissions the ACK frame follows the data transmission, while for uplink transmissions the acknowledgment is made using the ACK list of the next beacon frame. The piggybacking of acknowledgments in the

beacon frame eliminates the bandwidth overhead associated with the reception of individual ACK frames for each uplink data frame. As an example, in the 250 kbit/s version of the IEEE 802.15.4 standard, the transmission of an ACK frame (11 bytes) takes 352 μs and the gap between the data and ACK frames can vary from 192 to 512 μs. Even using the minimum gap, the overhead introduced by the use of ACK frames, considering the data frame length used in Section IV-C (43 bytes), would be 40%.

The ACK list is composed by an ACK length (AL) field and an ACK bitmap field containing one bit for each uplink RG of the previous superframe. A successful transmission is indicated by a "1" in the respective bitmap position, while a lost or corrupted transmission is indicated by a "0".

To reduce the power consumption of the stations, after a station listen to the beacon, it turns off the radio transceiver, turning it on again in the beginning of each scheduled data transmission or reception. After that, it turns off the transceiver again. The station turns on the transceiver again at the beginning of the next superframe to receive a new beacon, and so the cycle repeats. Each time the station receives a beacon it resynchronizes its clock with the base station.

C. Theoretical analysis

This section presents an analysis of the LPRT protocol with uplink data traffic, assuming an AWGN/BSC (Additive White Gaussian Noise/Binary Symmetric Channel) channel, which is later used for comparison with experimental data.

The packet error rate (PER_D) for an uplink data transmission is given by:

$$PER_D = 1 - (1 - BER)^{L_D}, \quad (1)$$

where L_D is the length of the data packet in bits and BER is the channel bit error rate. The packet error rate (PER_B) for a beacon transmission, with L_B as the beacon length in bits, is calculated in a similar way.

The delivery error rate (DER) is defined here as the probability of failure in the delivery of a data message from a given station to the base station. We consider that one data message (e.g., a sample from all sensors of a smart shirt) is generated at each superframe. For the LPRT protocol, the basic frame exchange transaction required for the delivery of an uplink data message consists of the transmission by the base station of a beacon frame at the start of a superframe followed by the transmission by a station of a data frame, at the position announced by the beacon, in the same superframe.

A successful transaction in this case requires the successful transmission of both the beacon and the data frame; therefore the delivery error rate in this case (DER_0) is:

$$DER_0 = 1 - (1 - PER_B)(1 - PER_D). \quad (2)$$

In order to increase the robustness of the communications, the LPRT protocol implements a contention-free retransmission mechanism that works by automatically rescheduling the transmission of a data message in the next superframe whenever the delivery in the current superframe fails. If the attempt of delivery of the data message fails in two consecutive superframes the message is dropped. The improved delivery error rate with one retransmission attempt is:

$$DER_1 = DER_0^{\,2}. \quad (3)$$

One can notice that the rescheduled transmission in the RP can be either a retransmission, if the data transmission in the previous superframe failed, or the first transmission of the data, if instead the beacon in the previously superframe was not received, since it precludes the corresponding data transmission. For the sake of simplicity, however, the rescheduled transmission will continue to be referred here as a retransmission.

The sensors samples included in the data message are taken just before its transmission, so the delay without retransmission is just the data transmission time, while the delay for a retransmitted data message is always less than one superframe period.

The average current consumption of the LPRT protocol without retransmissions (I_0) is given by:

$$I_0 = \frac{T_B + G_B + T_D + G_D}{T_{SF}}(I_{ON} - I_{OFF}) + I_{OFF} \quad (4)$$

where T_B and T_D are the beacon and data transmission times, G_B and G_D are the respective guard times, I_{ON} is the current consumption in the active state and I_{OFF} is the consumption in the sleep state. The guard times are the periods between the transition to the active state and either the effective transmission of data or the reception of the beacon.

The average current consumption with retransmissions increases by a factor that depends on the DER, and is given by:

$$I_1 = I_0 + \frac{T_D + G_D}{T_{SF}}(I_{ON} - I_{OFF})DER_0 \quad (5)$$

D. PC application

The TinyOS installation CD includes a java application called SerialForwarder that reads packets from the serial port and forwards those through a given TCP port to a network client application, and vice versa. Based on this java application, we developed an application that introduces several new functionalities.

The main window provides information about the state of the connection between the base station and each station in real-time (Figure 6), including the number of beacons not received by a station and the number of data frames not received by the base station, as well as the DER with and without retransmissions. The user interface presents this data both for the current values (evaluated from the last N samples) and cumulative values (since the start of data collection). The receiver signal strength indicator (RSSI)

and the battery voltage of the station are also presented. The data can also be registered in a file for posterior analysis.

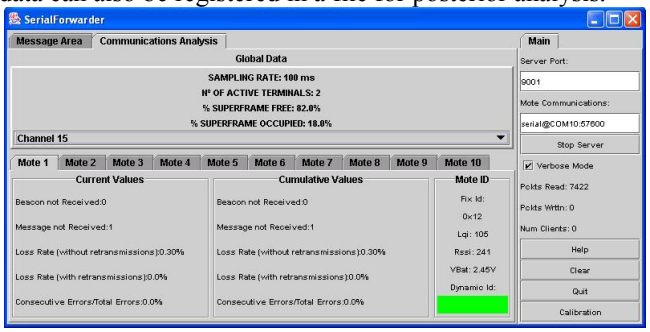

Figure 6. Main user interface of the PC application.

The developed PC application also includes a functionality that allows the sensors to be remotely calibrated though the wireless sensor network (Figure 7), thereby avoiding the requirement of physical connection of the data station to a programming board.

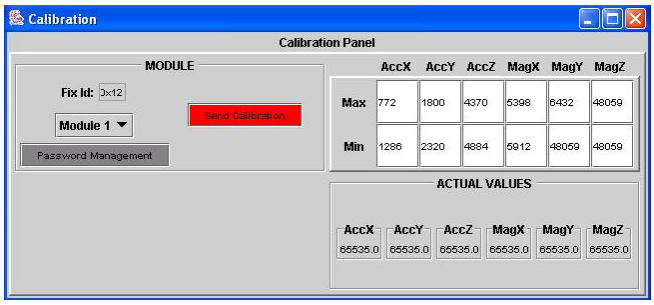

Figure 7. Calibration interface.

The application receives the raw posture sensors data through the wireless network, performs the calculation of the angles of the articulations and forwards this information to the client application, which acts as a data base and allows 3D visualization of data. Figure 8 presents an OpenGL application that was developed for real-time visualization of the posture and movements of the users.

E. *Experimental results*

First results in this section present the current consumption of the MICAz motes as a function of time, which include both the consumption of the transceiver and of the microcontroller. The evaluation scenario is composed by a base station and four stations. Each station collects 36 byte samples (corresponding to the data generated by four posture sensor modules) periodically, with a sampling rate of 5 Hz, which results in a packet inter-arrival time of 200 ms. The superframe period was chosen to be equal to the packet inter-arrival time, so one sample is transmitted per superframe. Figure 9 presents the current consumption for one of the stations. The current consumption of the microcontroller (I_{OFF}) is 8 mA while the consumption with the transceiver turned on (I_{ON}) is around 28 mA.

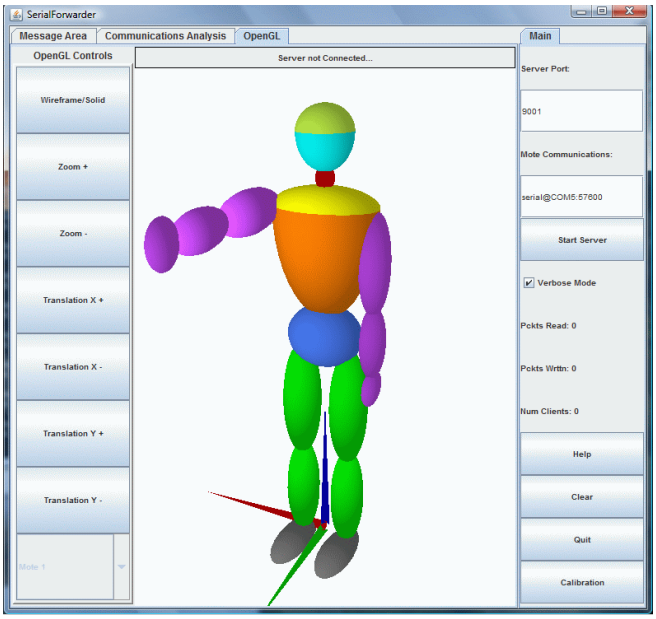

Figure 8. OpenGL application.

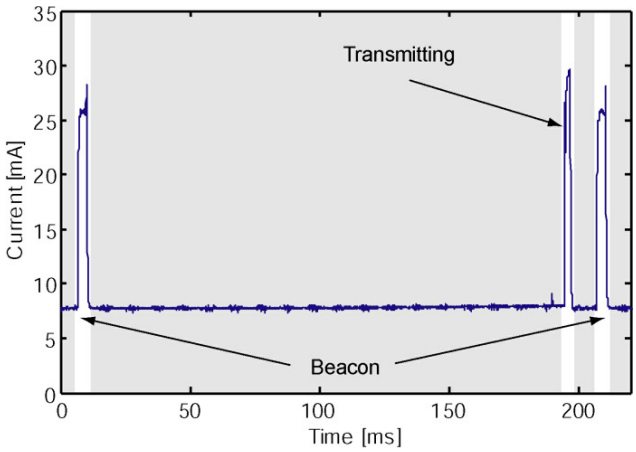

Figure 9. Current consumption of the LPRT protocol as a function of time for one particular station.

Figure 10 presents a closer look to the current consumption charts of the four stations in the RG list, superimposed, where the transmission of each station and the reception of the beacon of the next superframe can be seen. Each superframe is divided in 1024 mini-slots, allowing a high level of granularity in the allocation of transmission time to the stations. With these data frame size and without channel errors, the LPRT protocol allows up to 60 stations per superframe. It can achieve high throughput with high loads without losses due to collisions that are typical of contention based protocols.

The evaluation setup for the second experiment is composed by two networks operating in different channels (11 and 26) at the same time in the same place. Each network is composed by only one base station and one station located 3 meters apart and using the same

transmission power (0 dBm). The superframe period is T_{SF} = 100 ms. The samples are 75 bytes long, so the data frame length, including headers is L_D = 712 bits while the beacon length is L_B = 160 bits, which means that the respective transmission times, given that the network data rate is 250 kbit/s, are T_D = 2.85 ms and T_B = 0.64 ms.

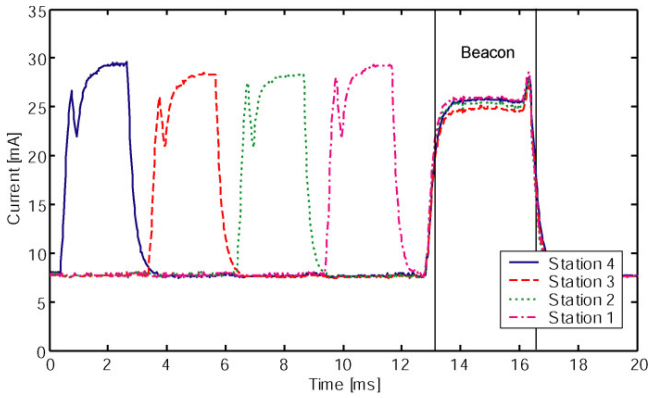

Figure 10. Current consumption of the LPRT protocol as a function of time for four stations.

Each station was powered by two 2300 mAh Ni-MH AA batteries. Both networks started to operate in a Thursday at 17:00h. The channel 11 network ceased to operate when the batteries drained out. The operation of channel 26 network lasted longer but was terminated due to a power outage in the laboratory.

MICAz is compliant with the 2.4 GHz version of the IEEE 802.15.4 standard at the physical layer, which provides 16 non overlapping channels (11 to 26) spaced 5 MHz apart. The center frequency for channel 11 is 2405 MHz, while for channel 26 the center frequency is 2480 MHz.

The signals of three IEEE 802.11 access points operating on channel 1 were detected at the place of the experiment. Channel 1 is centered on 2412 MHz and has a width of 22 MHz, so it overlaps with channel 11 of 802.15.4. 802.11 access points operating in other channels were also detected, but there were no overlapping with 802.15.4 channel 26.

Figure 11 presents the DER as a function of time calculated using the previous 1000 samples. Asterisks in the graph identify disconnections, triggered by several consecutive missing data frames, followed by new associations. The number of errors for the network operating in channel 11 was significantly higher (overall DER without retransmissions was 0.66% for channel 11 and 0.02% for channel 26). These results can be explained by interference from nearby 802.11 networks. In fact, the great majority of errors occurred during working hours.

In order to avoid interference, the user interface of the developed PC application allows the operator to change the operating channel (Figure 6). The information about the new channel is sent on the beacon, allowing a seamless channel transition for the stations. Since the state of the links is monitored in real-time by the application, an automatic channel selection mechanism could be easily implemented as well.

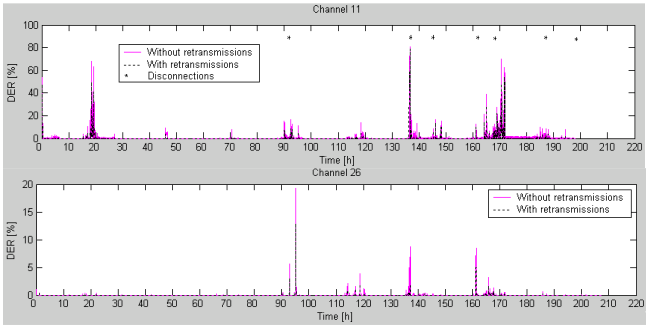

Figure 11. Delivery error rate as a function of time.

Figure 12 presents the relation between the DER with and without retransmissions, comparing the experimental data for channel 11 with the analytical data (equation 3). Even though there is a set of data points around the analytical curve, there is another set of points where the efficacy of the retransmission mechanism is almost none. In this situation, the overall performance is worse than expected, which indicates that future analysis needs to take into account other effects, such as burst errors due to fading and/or interference.

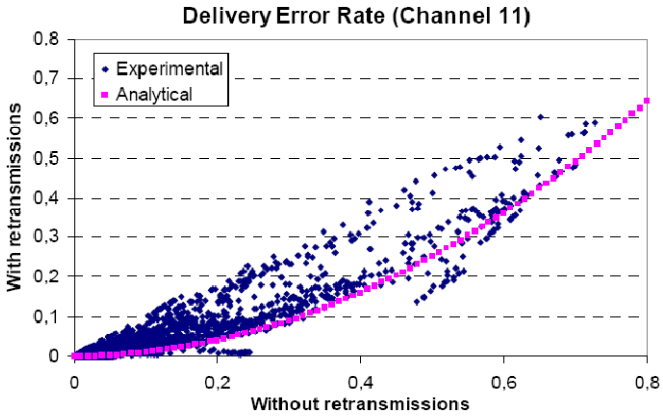

Figure 12. DER with and without retransmissions.

Figure 13 presents the battery voltage as function of time. It drops quickly when the voltage reaches around 2.2 V. The guard times that were used are G_D = 1 ms and G_B = 3.2 ms, thus, according to equation 4, the average current consumption without retransmissions is I_0 = 9.54 mA, so the theoretical lifetime of the station with fully charged 2300 mAh batteries is 241 h. An inspection of the channel 11 chart indicates that the lifetime of the channel 26 station would probably be around 230 h, which is close to the theoretical value, if it was not for the power outage. The lifetime of the station on channel 11 was lower, probably because its batteries were not with full capacity anymore, since the increment in consumption due to retransmissions (equation 5) is small. Even in the worst case (DER_0 = 1) it

would be only 8 %. The average current consumption of stations could be reduced to 2.15 mA if the microcontroller was turned off during the sleep state, increasing significantly the lifetime.

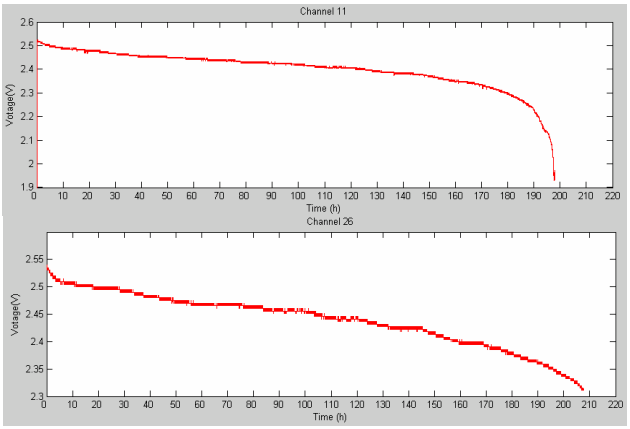

Figure 13. Battery voltage as function of time.

Another set of experiments was made in a hall with one base station and one station. A range (r) of 30 m with moderate channel errors (DER_0 < 10 %) was achieved when both the base station and the station were at the chest level. However, when the station was placed at ground level, the errors increased significantly and disconnections were frequent, even for shorter distances. This behavior can be explained by the radio signal power dropping off with r^4 [13] due to ground reflections for short antenna heights.

When the base station was placed over the station, the channel conditions were not affected by the height of the station antenna, even when the station was placed at the ground. This results lead to conclusion that a good place to put the base station is at the ceiling, in the middle of the hall, since it is a central position with relation to the stations and the signal is less affected by the height of the stations antennas.

IV. Second System Prototype

While the architecture of the first prototype comprises a wireless network of suits where the sensor nodes in each suit are connected through a wired BSN to a wireless station that communicates with a base station, the second prototype is based on a fully wireless architecture, i.e., all sensor nodes communicate wirelessly with the base station, as shown in Figure 14.

A. Prototype Overview

The architecture of this second prototype provides more flexibility of use compared to the first prototype, extending the range of applications of the system, since the number of sensor nodes and their distribution through the body can be easily changed according to the needs of the application.

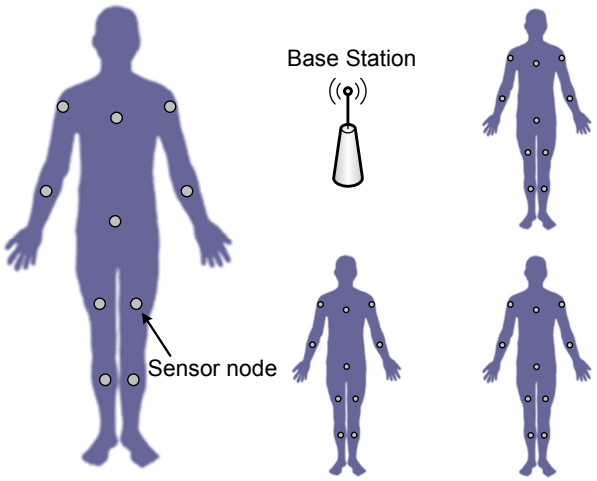

Figure 14. Architecture of the second prototype.

The design of the wireless sensor nodes requires a microcontroller, an analog-to-digital converter (ADC), a transceiver, an antenna and a battery, as well as the sensors. The sensor nodes should be as small and lightweight as possible, since the user may need to use several of them in the body.

The MICAz mote, used in the first prototype, is not adequate in this case, since its size, excluding the battery pack, is 58x32x7 mm. The second prototype is based on the CC2430 [14], a system-on-chip (SoC) from Texas Instruments that integrates a microcontroller, an ADC and a transceiver in a 7x7 mm QLP48 package. Like the MICAz mote, the CC2430 is compliant with the 2.4 GHz version of the IEEE 802.15.4 at the physical layer. Figure 15 presents a CC2430 evaluation module (EM) with external antenna. These modules are part of a development kit [15] used to program and debug the MAC protocol implemented in the second prototype.

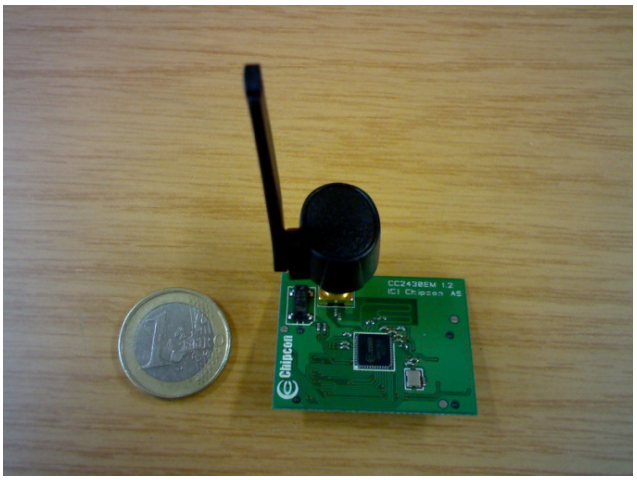

Figure 15. CC2430 evaluation module with external antenna.

The hardware of the prototype of the wireless sensor nodes consists of a PCB integrating a CC2430, an antenna, the sensors and the battery. Instead of using an external antenna, the sensor node will use a small size PCB antenna, described in [16].

B. MAC Protocol

The second prototype implements an improved version of LPRT, called iLPRT [17]. This protocol explores the fact that the traffic of BSN applications is mostly periodic to avoid the transmission of allocation information (RG list) in every superframe (a characteristic of LPRT). Instead of placing the RG list in the beacon, iLPRT includes only the CP size, which is necessary to avoid transmissions during the CFP from nodes that are using CSMA/CA. For transmissions during the NTP, nodes use the same allocation from the previous superframe. For transmissions in the RP, each node computes the period that is reserved for its retransmission with basis on the ACK bitmap information contained in the beacon, using a distributed algorithm that is shared by all nodes.

To ensure fairness in the allocation of retransmission periods when the available RP space is not sufficient to support all retransmissions, the algorithm performs rotation of allocation periods through all network nodes.

Unlike LPRT, if a node does not receive a beacon or a short sequence of beacons, it may continue to transmit new data in the NTP, since the clock drift should allow the node to continue synchronized during a few consecutive beacon intervals. Like LPRT, the node cannot retransmit in the RP when it does not receive the respective beacon, since the ACK bitmap is not available, so it does not know how the RP slots were allocated.

Since the RG list is not included in every beacon, a reconfiguration scheme, such as the algorithm proposed in [17], is required to inform the nodes when a new allocation is made or an existing allocation is released.

The use of shorter beacons tends to decrease the DER and the energy consumption of the sensor nodes. The possibility of transmission during the NTP when the beacon is lost also tends to decrease the DER. This is an advantage also when compared to the GTS mechanism of IEEE 802.15.4, since in this protocol if a device misses the beacon it must not use its allocated slots until it receives a beacon correctly.

C. Simulation Results

Simulation models of the LPRT and iLPRT protocols were implemented using the version 4.0 of OMNeT++, an open source, component-based, modular and open-architecture discrete event simulation environment [18] [19].

Table I presents the parameters related to the posture sensors and their corresponding values. Most motion capture applications operate with a frame rate of 30 fps (frames per second), so the sampling rate of the posture sensors (f_s) was set to 30 Hz. Since each posture sensor module has three accelerometers and three magnetometers, $N_s = 6$.

TABLE I. MOTION CAPTURE PARAMETERS.

Parameter	Name	Value
Number of sensors	N_s	6
Sampling rate	f_s	30 Hz
Sensors sampling resolution	Q_s	12 bits
Battery sampling resolution	Q_b	8 bits

Table II presents the network-related parameters and values. The PHY overhead corresponds to the physical preamble and header of IEEE 802.15.4. The MAC overhead is due to the MAC header and trailer. For data messages, it includes the following fields: FCF (Frame Control Field), data sequence number, PAN address, destination address, source AID, frame type and FCS (Frame Check Sequence).

TABLE II. NETWORK PARAMETERS.

Parameter	Name	Value
PHY data rate	R	250 kbit/s
PHY overhead	OH_{PHY}	6 bytes
MAC overhead	OH_{MAC}	9 bytes
Payload length	L_p	28 bytes
Superframe duration	T_{SF}	100 ms
Mini-slots per superframe	M_{SF}	500

The superframe duration (T_{SF}) was set to 100 ms. With this value, the number of samples (N_{sa}) of each sensor in the posture sensor module, using a sampling rate of 30 Hz, is 3, according to the equation:

$$N_{sa} = f_s T_{SF}. \quad (6)$$

Besides of the multiple samples of the posture sensors, each data message carries also a sample of the battery voltage. The length of the payload (L_p) of the data messages is given by the equation:

$$L_p = N_{sa} N_s Q_s + Q_b, \quad (7)$$

where N_s is the number of sensors, Q_s and Q_b are the sampling resolution of the posture sensors and the battery voltage, respectively. To save space, each two 12-bit samples are compressed into three bytes. With the values obtained from Table II, the payload length results in 224 bits, or 28 bytes.

There is a tradeoff between the delay and the energy consumption and network capacity (number of supported nodes). An increase in the superframe duration would tend to increase the network capacity and decrease the energy consumption, since the payload length would increase, reducing the protocol overhead. However, the delay would also tend to increase.

The maximum PPDU (PHY Protocol Data Unit) size of 133 bytes, imposed by the IEEE 802.15.4 specifications, limit the number of RGs in the beacon and, consequently,

the number of transmissions in the CFP for LPRT, to a maximum of 55. On the other hand, iLPRT does not suffer this constraint since the RG list is not included in the beacon.

Comparing with the RG field format used in the first prototype, the length of the AID field was increased from 5 to 6 bits, while the length of the ITS field was decreased from 10 to 9 bits, which means that the RG field maintains the overall length of 2 bytes. With 6 bits in the AID, the addressing capability of the network is increased to 64 nodes. With 9 bits in the ITS field, the superframe can have up to 512 mini-slots. The number of mini-slots per superframe (M_{SF}) was set to 500 instead, because with this value the period of each mini-slot, with the superframe duration of 100 ms, is 200 µs, which is an integer multiple of the resolution of the timer of the CC2430 that is being used (4 µs).

The minimum duration of the CP (CP_{min}) was set to 11 ms, which is enough to support the exchange of two full length frames, including the CSMA/CA backoff and the gap between the frames. The CP is used, for example, for the exchange of association request/response messages between the nodes and the base station and for the transmission of sporadic sensor calibration data. The maximum duration of the CFP (CFP_{max}) is the superframe duration (T_{SF}) minus a reserved period (15.26 ms) composed by the CP_{min} plus the time required to transmit a maximum length beacon (4.26 ms). Therefore, given the used superframe duration, CFP_{max} = 84.74 ms, which corresponds to 423 mini-slots. According to the values defined on Table II, the length of the PPDU (Physical layer Protocol Data Unit) is 43 bytes. That means that the transmission time is 1.376 ms, requiring 7 mini-slots. One additional mini-slot was used as guard time between the transmissions of the sensor nodes, resulting in 8 reserved mini-slots for each data message. Consequently, the maximum number of transmissions that can be carried out in the CFP is 52, either with LPRT or iLPRT.

A brief analysis of the IEEE 802.15.4 GTS mechanism shows that the limited options for the superframe duration combined with the low granularity of the slots result in poor bandwidth efficiency. The IEEE 802.15.4 standard provides only 15 choices for the superframe duration, ranging in a geometric progression from 15.36 ms to 251.658 s. The closest options to the value defined in Table II (100 ms) are 61.44 ms and 122.88 ms. Using the latter value, the average number of samples per superframe with 30 fps would be 3.69 and the average payload length (applying equation 7) would be 34.21 bytes, so the data frame transmission time would be 1.58 ms. Since each superframe is composed by 16 slots, each slot would have a period of 7.68 ms, which means that 79.5% of the available slot space would be wasted if the GTS mechanism was used.

The next three figures present the simulation results for the delivery error rate (DER) as a function of the number of sensor nodes (considering different channel error models and parameter values). For each figure the following cases were simulated:

- LPRT without retransmissions;
- LPRT with one retransmission attempt;
- iLPRT without retransmissions;
- iLPRT with one retransmission attempt.

Simulation results presented in Figure 16 assume an AWGN/BSC channel identical to the used in the theoretical analysis presented on Section III-C. It is also considered that the channel is symmetrical in terms of bit error rate, which means that the BER used in the transmission of the beacon by the base station is the same BER used in the transmission of the data messages by the sensor nodes.

The simulation results for the LPRT protocol are identical to the ones obtained with the theoretical analysis, using equations 2 and 3, for the cases without retransmissions and with one retransmission attempt, respectively. The consistency of results between these two analysis methods provides validation to the simulation model of the LPRT protocol.

As shown by the figure, the DER with LPRT increases with the number of the nodes due to the increase of the beacon length. This is mainly due to the corresponding increase in the RG list length, which takes two bytes per allocated transmission (Figure 5). The beacon of iLPRT does not carry the RG list and its length only increases by one bit per node due to the ACK bitmap, so the DER with iLPRT is practically independent of the number of nodes in the network.

The DER with retransmissions is much better than the DER without retransmissions, for both protocols, when there is sufficient space available in the superframe for all retransmissions. As the number of nodes approaches the capacity of the network, the number of scheduled retransmissions decreases, due to lack of space in the superframe, increasing the DER. For example, the DER with iLPRT starts to increase with more than 48 nodes. In the limit, when the number of nodes reaches 52 and there is no space for retransmissions anymore, the DER with the retransmission mechanism for both protocols reaches the corresponding value of the DER without retransmissions.

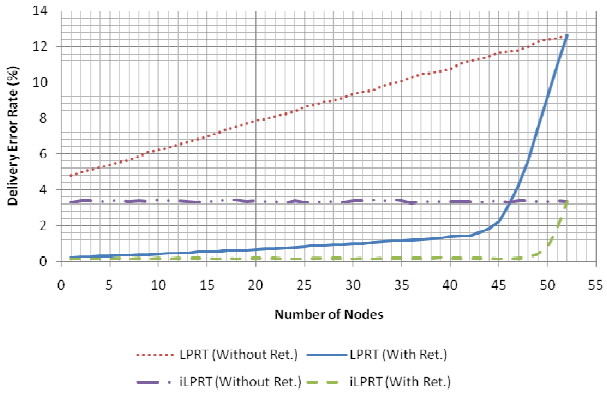

Figure 16. DER using the AWGN/BSC model with BER = 10^{-4}.

As the figure shows, iLPRT improves the DER with respect to LPRT. The main reasons for this improvement are the shorter beacons used by iLPRT, that are less sensitive to channel errors, and the transmission of data by the nodes during the NTP when the beacon is not received, which is allowed with iLPRT but not with LPRT.

A second scenario that takes into account the occurrence of burst errors was considered. The channel errors for this scenario were modeled using the Gilbert-Elliot model [20]. According to this model, the channel alternates between a good state with low bit error rate (BER_{good}) and a bad state, with very high bit error rate (BER_{bad}), with mean dwelling time T_{good} for the good state and T_{bad} for the bad state. The choice of values for these parameters, presented in Table III, is intended to model fast fading, which typically occurs on timescales of milliseconds to tens of milliseconds [21] and can also represent interference from other sources such as IEEE 802.11 transmissions in the same band. The channel errors for the different nodes were made independent, which means that at any moment the channel for some nodes can be in the bad state while for others it can be in the good state. The channel state is the same regardless of the direction of the transmission (from the node to the base station or vice-versa).

TABLE III. PARAMETERS OF THE GILBERT-ELLIOT MODEL.

Parameter	Value
BER_{bad} (nodes to base station)	10^{-2}
BER_{bad} (base station to nodes)	$10^{-2}/10^{-3}$
BER_{good}	0
T_{bad}	10 ms
T_{good}	90 ms

Figure 17 presents the DER curves obtained using the same value of BER_{bad} (10^{-2}) in both directions. The average BER, using the parameters of Table III is 10^{-3}, which is significantly higher than the constant BER used in the previous scenario.

With this value of BER_{bad}, the probability that a beacon or a data message is corrupted by errors when the channel is in the bad state is near 100%, which means that almost all successful transmissions occur only when the channel is in the good state. Since in this scenario all beacons tend to be corrupted when the channel is in the bad state, regardless of their length, the DER using LPRT without retransmissions shows little sensitivity to the number of nodes in the network, unlike what occurs in the previous scenario (Figure 16).

Like in the previous scenario, iLPRT presents an improvement in the DER with respect to LPRT; however, the DER observed in the figure, for both protocols, is worse with small rather than medium number of nodes in the network. The explanation for this behavior relates to the burst errors channel model and the way the transmissions are allocated in the superframe. The scheduling of mini-slots for transmission of data messages starts from the end of the superframe and goes towards to its beginning, as the number of nodes increases, which means that the transmission of the first node is closely followed by the beacon of the next superframe. In a scenario with burst errors, the probability that the bad state observed during the transmission of this data message extends into the reception of the beacon is high. When the beacon is corrupted, the node is unable to make the retransmissions, since it does not have the feedback about the delivery of the data message neither the information about the allocation of mini-slots for retransmission. Therefore, burst errors tend to reduce the effectiveness of the retransmission mechanism for data messages that are scheduled for transmission close to the next beacon. As the number of nodes in the network increases, the average distance between the data transmissions and the following beacon also increases, decreasing the probability that the reception of the beacon is compromised by the bad state observed during the transmission of the data message.

The behavior observed in Figure 17 provides an explanation for the experimental results of DER with retransmissions using LPRT that was presented in Figure 12 (in a scenario with a single sensor node and equal output power in both directions). The experimental results are worse than expected by the theoretical results provided by equation 3, which consider a channel with constant BER, while the simulation results obtained in Figure 17 are closer. This suggests that, in similar circumstances, channel models that take into account burst errors tend to be more realistic than models that assume an AWGN/BSC channel.

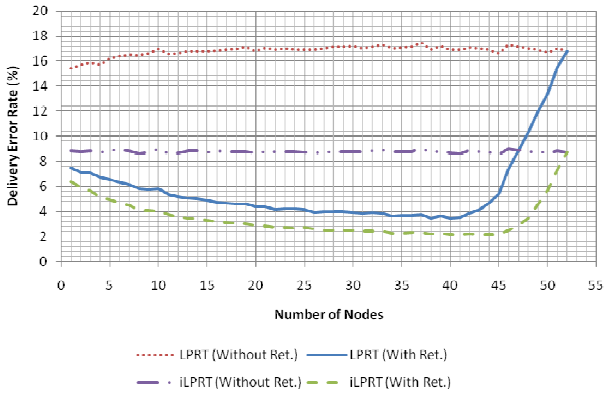

Figure 17. DER using the Gilbert-Elliot model with $BER_{bad} = 10^{-2}$.

The BER at the receiver depends, among other factors, of the modulation used and the SNR (signal-to-noise ratio) [22]. Since the SNR depends of the output power, a solution to improve the robustness of these protocols consists in increasing the output power of the base station transmitter. This increase is not problematic in terms of energy consumption since, unlike the sensor nodes, the base station is not energy constrained.

Figure 18 presents the DER curves obtained in a scenario similar to the previous one. The only difference is

that the BER_{bad} for the transmissions of the beacon from the base station to the nodes is reduced to 10^{-3}, to account for an increase in the base station output power, while the BER_{bad} for the transmissions from the node to the base station maintains the value of 10^{-2}. As Figure 18 shows, the increment in the probability of reception of the beacon when the channel is in the bad state has a overall positive effect on the DER with retransmissions for both protocols, but especially when the number of nodes is small, correcting the problem detected in the previous scenario.

Since the BER_{bad} for the beacon in this scenario is lower than in the previous one, the probability of corruption of the beacon becomes significantly dependent of the beacon length again and, consequently, the DER using LPRT without retransmissions shows a dependence on the number of nodes in the network. On the other hand, the DER using iLPRT without retransmissions in this scenario remains the same when compared to the previous scenario, even though the probability of corruption of the beacon is different, since the transmission of data messages during the NTP with iLPRT does not depend of the successful reception of the beacon.

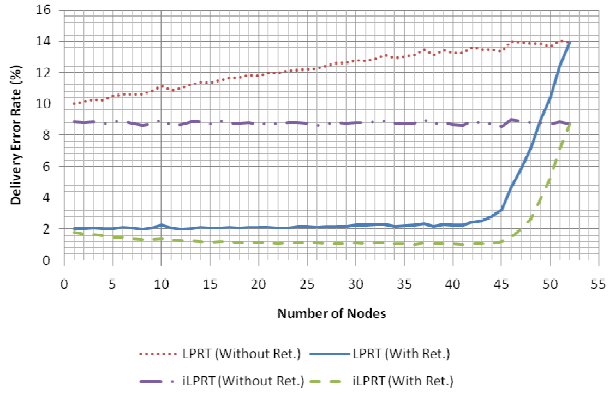

Figure 18. DER using the Gilbert-Elliot model with $BER_{bad} = 10^{-2}/10^{-3}$.

Figure 19 shows the average current consumption (per node) of the CC2430 circuitry with the retransmission mechanism enabled, considering the current consumption of 26.7 mA (RX mode) and 26.9 mA (TX mode, 0 dBm) in the active state and 0.5 µA (power mode 2) in the sleep state. The consumption using LPRT increases significantly with the number of nodes, mainly because of the increment in the size of the beacon due to the RG fields (the consumption starts to decrease with more than 45 nodes because the nodes start to drop the retransmissions due to lack of space in the superframe). The consumption of iLPRT increases much less, because the increment on the size of the beacon in this case is only due to ACK bitmap field. The increase in the current consumption due to retransmissions is small, being largely compensated by the benefit of reduction of the DER.

The average current consumption of the CC2430 for the operation of the iLPRT protocol with the retransmission mechanism enabled, considering a scenario of utilization with 45 sensor nodes, is 0.6 mA. The ADC sampling time for each sensor is 132 µs, which means that the total sampling time for the six sensors in the posture sensor module is 792 µs. With a sampling rate of 30 Hz, that results in a duty-cycle for the ADC of 2.38%. Since the consumption of the microcontroller during the ADC operation is 12.3 mA, the average current consumption due to the ADC using this duty-cycle is 0.29 mA.

Most of the consumption of the sensor node is due to the posture sensors circuitry, which consume around 11.5 mA. The total current consumption including the operation of the iLPRT protocol and the ADC is therefore 12.39 mA. Considering the use of a rechargeable Lithium-Ion battery with capacity of 300 mAh, the lifetime of each posture sensor module will be in the order of 24.2 h of continuous operation. The lifetime of the nodes could be significantly increased by switching off part of the posture sensors circuitry between samples.

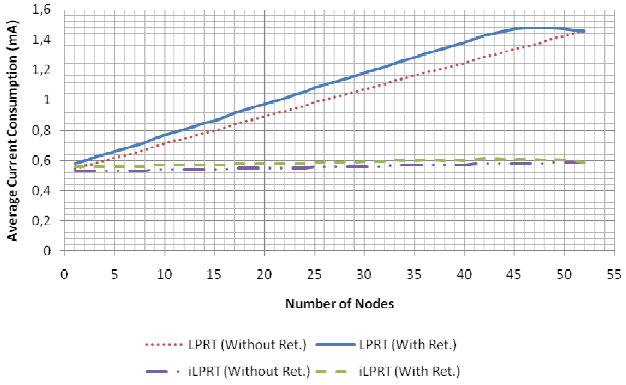

Figure 19. Current consumption with $BER_{bad} = 10^{-2}/10^{-3}$.

V. Conclusions and Future Work

This paper describes the design and development of two multi-user wireless BSN prototypes where the posture and movements of multiple users can be monitored simultaneously in real-time using a single base station. The first prototype is implemented using MICAz motes and the second one is based on the CC2430 integrated circuit.

The developed prototypes use the physical layer of the IEEE 802.15.4 protocol, but replace the MAC layer with LPRT/iLPRT. The main advantages of these protocols over the IEEE 802.15.4 CSMA/CA protocol are the collision-free operation and the immunity to the hidden node problem. Advantages of LPRT/iLPRT over the GTS mechanism are: support for more than seven nodes; higher granularity in the allocation of transmission periods; piggybacking of the ACK feedback in beacon frames and automatic scheduling of collision-free retransmissions. These advantages tend to increase the bandwidth efficiency. Another advantage of iLPRT over the GTS mechanism is that the nodes can make the normal data transmission in the superframe even if the

beacon is not received, which tends to decrease the delivery error rate.

Future work includes the porting of the base station code of the second prototype from the CC2430 to the CC2530 [23] integrated circuit, which can supply more 4.5 dB of output power than the CC2430, providing a simple alternative to increase the probability of reception of the beacon by the sensor nodes and, consequently, enhance the performance of the system. After that, a set of laboratory and field tests will be performed to evaluate the performance of the system in different scenarios.

ACKNOWLEDGEMENT

This work is supported by Portuguese Foundation for Science and Technology (FCT) project PTDC/EEA-TEL/68625/2006.

REFERENCES

[1] J. A. Afonso, H. R. Silva, P. M. Oliveira, J. H. Correia, and L.A. Rocha, "Design and Implementation of a Real-Time Wireless Sensor Network", *International Conference on Sensor Technologies and Applications (SENSORCOMM 2007)*, Valencia, Spain, October 2007.

[2] B. Lo and G. Z. Yang, "Key Technical Challenges and Current Implementations of Body Sensor Networks", *2nd International Workshop on Body Sensor Networks (BSN 2005)*, London, UK, April 2005.

[3] M. Paksuniemi, H. Sorvoja, E. Alasaarela, and R. Myllylä, "Wireless Sensor and Data Transmission Needs and Technologies for Patient Monitoring in the Operating Room and Intensive Care Unit", *27th Annual International Conference of the Engineering in Medicine and Biology Society (IEEE EMBC 2005)*, Shanghai, China, September 2005.

[4] IEEE Std 802.11-2007, IEEE Standard for Information Technology—Telecommunications and information exchange between systems—Local and metropolitan area networks—Specific requirements—Part11: Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications, June 2007.

[5] Bluetooth SIG, "Specification of the Bluetooth system", November 2003. Retrieved from http://www.bluetooth.org.

[6] IEEE Std 802.15.4-2006, IEEE Standard for Information Technology—Telecommunications and information exchange between systems—Local and metropolitan area networks—Specific requirements—Part 15.4: Wireless Medium Access Control (MAC) and Physical Layer (PHY) Specifications for Low-Rate Wireless Personal Area Networks (WPANs), September 2006.

[7] ZigBee Standards Organization, "ZigBee Specification", 2006. Retrieved from http://www.zigbee.org.

[8] D. Marculescu, R. Marculescu, S. Park, and S. Jayaraman, "Ready to Ware", *IEEE Spectrum*, Vol. 40, Iss. 10, pp. 29-32, October 2003.

[9] M. Catrysse et al., "Towards the Integration of Textile Sensors in a Wireless Monitoring Suit", *Sensors and Actuations A: Physical*, Vol. 114, Iss: 2-3, pp. 302-311, September 2004.

[10] J. A. Afonso, J. H. Correia, H. R. Silva, and L. A. Rocha, "Body Kinetics Monitoring System", International Patent WO/2008/018810A2 , February 2008.

[11] Crossbow Technology Inc, "MPR/MIB User's Manual", April 2005. Retrieved from http://www.xbow.com.

[12] D. Gay et al., "The nesC Language: A Holistic Approach to Networked Embedded Systems", *ACM SIGPLAN 2003 conference on Programming language design and implementation*, pp. 1-11, San Diego, CA, June 2003.

[13] D. Estrin, L. Girod, G. Pottie, and M. Srivastava, "Instrumenting the World with Wireless Sensor Networks", *IEEE International Conference on Acoustics, Speech, and Signal Processing (ICASSP '01)*, Salt Lake City, UT, May 2001.

[14] Texas Instruments, "CC2430 Data Sheet (rev. 2.1)", 2007. Retrieved from http://www.ti.com.

[15] Texas Instruments, "CC2431DK Development Kit User Manual Rev. 1.5", 2007. Retrieved from http://www.ti.com.

[16] Audun Andersen, "Small Size 2.4 GHz PCB antenna", Application Note AN043, 2008. Retrieved from http://www.ti.com.

[17] O. Gama, J. A. Afonso, P. Carvalho, and P. M. Mendes, "An Improved MAC Protocol with Reconfiguration Scheme for Wireless e-Health Systems Requiring Quality of Service", *Wireless VITAE 2009*, Aalborg, Denmark, May 2009.

[18] A. Varga, "The OMNeT++ Discrete Event Simulation System" *European Simulation Multiconference (ESM 2001)*, Prague, Czech Republic, June 2001.

[19] OMNeT++ Community Site. http://www.omnetpp.org.

[20] J. Ebert and A. Willig, "A Gilbert-Elliot Bit Error Model and the Efficient Use in Packet Level Simulation", *TKN Technical Report TKN-99-002*, March 1999.

[21] A. Willig, "Recent and Emerging Topics in Wireless Industrial Communications: A Selection*", IEEE Transactions in Industrial Informatics*, Vol. 4, No. 2, May 2008.

[22] M. Zuniga and B. Krishnamachari , "Analyzing the Transitional Region in Low Power Wireless Links", *1st IEEE Annual Conference on Sensor and Ad Hoc Communications and Networks (SECON 2004)*, pp. 517-526, Santa Clara, CA, October 2004.

[23] Texas Instruments, "CC2530 Data Sheet", April 2009. Retrieved from http://www.ti.com.