



Technical Report

BANMAC: An Opportunistic MAC Protocol for Reliable Communications in Body Area Networks

Shashi Prabh

Fernando Royo

Stefano Tennina

Teresa Olivares

HURRAY-TR-120401

Version:

Date: 4/8/2012

BANMAC: An Opportunistic MAC Protocol for Reliable Communications in Body Area Networks

Shashi Prabh, Fernando Royo, Stefano Tennina, Teresa Olivares

IPP-HURRAY!

Polytechnic Institute of Porto (ISEP-IPP)

Rua Dr. António Bernardino de Almeida, 431

4200-072 Porto

Portugal

Tel.: +351.22.8340509, Fax: +351.22.8340509

E-mail: ksph@isep.ipp.pt, , sota@isep.ipp.pt,

<http://www.hurray.isep.ipp.pt>

Abstract

We consider reliable communications in Body Area Networks (BAN), where a set of nodes placed on human body are connected using wireless links. In order to keep the Specific Absorption Rate (SAR) as low as possible for health safety reasons, these networks operate in low transmit power regime, which however, is known to be error prone. It has been observed that the fluctuations of the Received Signal Strength (RSS) at the nodes of a BAN on a moving person show certain regularities and that the magnitude of these fluctuations are significant (5 - 20 dB). In this paper, we present BANMAC, a MAC protocol that monitors and predicts the channel fluctuations and schedules transmissions opportunistically when the RSS is likely to be higher. The MAC protocol is capable of providing differentiated service and resolves co-channel interference in the event of multiple co-located BANs in a vicinity. We report the design and implementation details of BANMAC integrated with the IEEE 802.15.4 protocol stack. We present experimental data which show that the packet loss rate (PLR) of BANMAC is significantly lower as compared to that of the IEEE 802.15.4 MAC. For comparable PLR, the power consumption of BANMAC is also significantly lower than that of the IEEE 802.15.4. For co-located networks, the convergence time to find a conflict-free channel allocation was approximately 1 s for the centralized coordination mechanism and was approximately 4 s for the distributed coordination mechanism.

BANMAC: An Opportunistic MAC Protocol for Reliable Communications in Body Area Networks

K. Shashi Prabh*, Fernando Royo*[†], Stefano Tennina*, Teresa Olivares[†]

* Center for Real-Time Systems Research, School of Engineering, Polytechnic Institute of Porto, Portugal

ksp@isep.ipp.pt, sota@isep.ipp.pt

[†] Albacete Research Institute of Informatics, University of Castilla La Mancha, Spain

froyo@dsi.uclm.es, teresa@dsi.uclm.es

Abstract—We consider reliable communications in Body Area Networks (BAN), where a set of nodes placed on human body are connected using wireless links. In order to keep the Specific Absorption Rate (SAR) as low as possible for health safety reasons, these networks operate in low transmit power regime, which however, is known to be error prone. It has been observed that the fluctuations of the Received Signal Strength (RSS) at the nodes of a BAN on a moving person show certain regularities and that the magnitude of these fluctuations are significant (5 – 20 dB). In this paper, we present BANMAC, a MAC protocol that monitors and predicts the channel fluctuations and schedules transmissions opportunistically when the RSS is likely to be higher. The MAC protocol is capable of providing differentiated service and resolves co-channel interference in the event of multiple co-located BANs in a vicinity. We report the design and implementation details of BANMAC integrated with the IEEE 802.15.4 protocol stack. We present experimental data which show that the packet loss rate (PLR) of BANMAC is significantly lower as compared to that of the IEEE 802.15.4 MAC. For comparable PLR, the power consumption of BANMAC is also significantly lower than that of the IEEE 802.15.4. For co-located networks, the convergence time to find a conflict-free channel allocation was approximately 1 s for the centralized coordination mechanism and was approximately 4 s for the distributed coordination mechanism.

I. INTRODUCTION

Due to the trend of a steadily increasing proportion of elderly people [1], the pressure on health-care services will continue to increase in the coming years. However, it is projected that there will be a severe shortage of nurses in the near future – the deficit is predicted to be about 1 million by the year 2020 in the USA alone [2]. Body Area Networks (BAN) will play a significant role in dealing with this shortage since they are capable of automating recording of the vital statistics, a job that consumes about forty percent of the nurses’ time. In addition, among other things, these networks facilitate better diagnosis, fast emergency response and personalized medication [3].

In the interest of protecting the human tissues, it is desirable that the transmission power be kept low. Low-power transmissions are known to suffer from aggravated communication errors. However, health-care applications generally require high reliability, since the cost of failure can be damage or loss of life of a person. In a previous work, we reported an empirical characterization of the Received Signal Strength Indicator (RSSI) fluctuations in BANs on walking people using

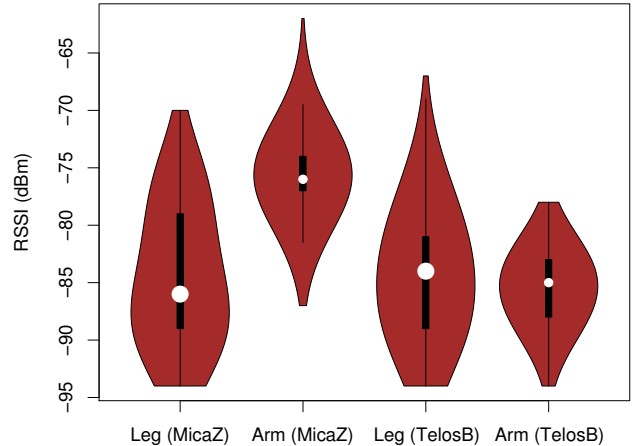


Fig. 1: Typical RSSI fluctuations on TelosB and MicaZ platforms. [4]

several hardware platforms [4]. In one set of the experiments, the coordinator was fixed at the chest and the other nodes were placed on lower leg and upper arm. Fig. 1, which is representative of the results, shows violin-plots of the RSSI fluctuations. The inter-quartile ranges (IQR) are shown as thin black boxes and the median is shown with a white dot. As shown in the figure, we found an IQR link margin of approx. 5 dB for the chest-arm pairs and of approx. 10 dB for chest-leg pairs. In another set of experiments using different set-up, we found RSSI fluctuations up-to approx. 20 dB. For more details, the reader is referred to the paper cited above.

The Received Signal Strength (RSS) fluctuations are the result of motion of the person wearing the network. The signal strength is higher when the nodes are close to line of sight and weaker when the body shadowing is strong. These channel fluctuations can be used advantageously since by appropriately timing the transmissions, we are more likely to get higher signal strength than average, which affects the communication reliability positively. We use the term *Opportune Transmission Window* (OTW) to describe a time interval that yields high RSS values relative to the average RSS of the link.

Normal human movements, such as walking, jogging or running, result in *approximately* periodic behavior of RSS fluctuations. During the time interval when the periodicity is sustained, it is possible to predict the occurrence of OTWs

in future. Predicting the next occurrences of OTWs requires dealing with the issues of noisy RSSI data and irregularity of human movements. In [4], we presented an RSSI-based OTW prediction algorithm and demonstrated the feasibility of this approach by evaluating the algorithm on real-world data. We also presented a preliminary sketch of leveraging these ideas to design an opportunistic MAC protocol, which we named BANMAC. However, we did not implement BANMAC in [4].

ZigBee provides a set of globally accepted specifications for wireless sensor networks. ZigBee standards span a wide range of applications including health, wellness and fitness [5]. ZigBee Health Care working group addresses medical care for the aging population, general wellness, sports training, etc. [6]. The MAC protocol used in ZigBee specifications is the IEEE 802.15.4 MAC and this is one of the motivations behind implementing BANMAC integrated with the IEEE 802.15.4 protocol stack.

The main contributions of the work presented in this paper are as follows:

- System design and implementation of BANMAC, a collision-free MAC protocol for body area networks. In BANMAC, the data transmissions are *scheduled*, as opposed to random channel access. BANMAC detects whether a node is on a mobile limb. For nodes on mobile limbs, it predicts the center of OTWs when the RSS of the transmissions of the mobile nodes is likely to be higher, which leads to higher reliability. Transmissions from stationary nodes can be scheduled during the rest of the available time. We note that BANMAC is flexible with respect to applying user defined scheduling policy.
- Support for multiple co-located BANs. BANMAC features both centralized and fully distributed coordination mechanisms. It uses the services of global coordinator whenever available and seamlessly switches to fully distributed mode once the network goes out of the range of the coordinator.
- Extensive evaluation of the opportunistic transmission mechanism using BANMAC. We compare the performance of BANMAC with that of IEEE 802.15.4 MAC with CSMA/CA. We also evaluate the capability of BANMAC to provide differentiated service.

The rest of this paper is organized as follows: in Sec. II we present network model and also a brief summary of our previous work to make the presentation self-contained. In Sec. III, we present the design and implementation details of BANMAC. We follow this by presenting the results of experimental evaluations in Sec. IV. We discuss related work in Sec. V and conclude the paper in Sec. VI.

II. PRELIMINARIES AND BACKGROUND

In this section we outline the network model and a brief summary of the OTW prediction method presented in [4] which should be useful for following the rest of this paper.

Network Model

BANs are typically arranged in a star topology where a set of nodes are wirelessly connected to a (BAN) coordinator. The coordinator can be connected to external networks. Generally, powerful devices like cellphones or PDAs are well suited for a coordinator. The nodes, however, are assumed to have limited energy supply and limited processing power. The coordinator is typically significantly more powerful than the rest of the nodes. Therefore, it is desirable to push as much computation and communication overhead to the coordinator as possible (BANMAC is indeed pull-based).

RSSI-based OTW Prediction Algorithm [4]

The OTW prediction algorithm for a node essentially works in a cycle of three steps.

In the first step, the coordinator collects RSSI time-series as follows. It periodically broadcasts RSSI probe packets at sufficiently low frequency, which are usually interspersed between data packets. Nodes record the RSSI of these probe packets and store them locally until the coordinator requests this data. The set of RSSI values with the probe identifiers (sequence number or coordinator's time-stamp) is sent back to the coordinator, aggregated in one or more packets. If the size of the time-series is small enough, it can also be piggybacked on data packets.

In the second step, the coordinator processes the (noisy) RSSI time-series to determine the frequency and phase of the RSSI fluctuation (due to the human movements). The algorithm achieves this by first applying Fast Fourier Transform (FFT) to the RSSI time-series and identifying the dominant frequency. As expected, our experiments confirmed that the dominant frequency in the FFT spectrum corresponds to the step frequency of the subject and the frequency components of noise tend to have much smaller amplitude. To determine the phase, we first apply a tight bandpass filter centered at the dominant frequency and subsequently apply an extrema identification algorithm.¹ Figure 2 shows the filtered signal superimposed on a sample of raw RSSI time-series.

In the third step, the algorithm predicts OTWs using the frequency and phase information. Clearly, the OTW should be centered at integral multiples of the period ($1/\text{dominant frequency}$) from an RSSI fluctuation maxima. Note that since the phases of the RSSI time-series in the beginning as well as in the end are arbitrary, we select the last but one peak from the previous step of extrema identification as the basis for OTW predictions.

Since normal human movement is irregular, we re-run the algorithm intermittently. For example, in our current beacon-enabled mode implementation of BANMAC where the beacons of 802.15.4 MAC are also used as RSSI probe packets, we run the OTW prediction algorithm every 64 beacon intervals. The repeating interval of these three steps can be chosen adaptively by monitoring the packet loss rate. We also note

¹Bandpass filters introduce additional phase. Alternatively, one can use forward-backward bandpass filter with the same results.

that since the limbs of humans move in synchrony – right hand and left leg move together and so do left hand and right leg – the OTWs for only one node on these moving limbs are sufficient to infer the OTWs for all other nodes. The OTWs for nodes on the right hand and left leg and those for nodes on the left hand and right leg alternate every half period (see Fig. 7).




Fig. 2: Raw and filtered RSSI time-series (TelosB node).

We evaluated the accuracy of our OTW predictions algorithm on real-world data collected on three different platforms (MicaZ, TelosB and Shimmer2) using two different scenarios. Figure 3 shows the absolute deviations between the center of predicted OTWs and the nearest bandpass filtered RSSI peaks observed in the experiments (or, “drifts”). Out of an RSSI time-series collected at 20 Hz, we used samples of 4.5 s duration every 12 s to predict the OTWs. The sampling time of 4.5 s includes at-least one pair of consecutive RSSI peaks. In the figure, the drifts less than $0.25 * period$ are shown with diamonds, those less than $0.5 * period$ and greater than $0.25 * period$ with triangles and the larger drifts with circles. The mean values of the absolute deviations were 0.28 s for the MicaZ samples, 0.21 s for the TelosB samples and 0.18 s for the Shimmer2 samples. The circles and triangles are sometimes at the same level as diamonds because of the variations in walking speed of the subject. This algorithm is robust in the event of moderate ($\approx 20\%$) packet losses. To ensure the exact periodicity of the samples for FFT, substituting the last reported RSSI for a lost data results in better accuracy than either setting RSSI to the minimum or maximum [7].

III. IEEE 802.15.4 BASED BANMAC

In this section, we present the system design of BANMAC integrated with the IEEE 802.15.4 MAC protocol, the most commonly used MAC in wireless sensor networks. As we mentioned previously, ZigBee standards span a wide range of applications including health, wellness and fitness [6]. The MAC protocol used in ZigBee specifications is also




Fig. 3: Accuracy of the OTW predictions.

IEEE 802.15.4. In the following, we first provide a brief overview of the relevant features of the IEEE 802.15.4 MAC, and follow it with the details of BANMAC design and implementation, where we present the cases of single network and co-located networks separately.

The IEEE 802.15.4 MAC standard defines two modes of operation, namely, beacon-enabled mode and nonbeacon-enabled mode [8]. In the beacon-enabled mode, the duration between the starting times of consecutive beacons, called Beacon Interval (BI), defines a TDMA “superframe”. One of the goals of the standard is low-power operation. To this end, only a part of the superframe is used for transmissions and during the rest of it, called Inactive Period, nodes go into sleep mode. The lengths of superframes and active periods are determined by two parameters, Beacon Order (BO) and Superframe Order (SO), respectively (see Fig. 4). Each active period of a superframe has 16 time slots. These slots are divided into Contention Access Period (CAP) where, within each slot nodes transmit using (slotted) CSMA/CA, and Contention Free Period (CFP) where, exclusive transmission rights are allocated to nodes. Within a CAP slot, more than one transmission can take place. In addition to specifying the superframe structure, beacons facilitate node synchronization and PAN identification. In the nonbeacon-enabled mode, the transmissions of a PAN coordinator to the end devices can take place directly only for broadcasts. Unicasts are done only indirectly: unicast packets from the coordinator are delivered only after the node has requested data from the coordinator. In the reverse direction, however, nodes can send packets to the coordinator using (unslotted) CSMA/CA directly.

We use the reserved bits 7–9 of the frame control field for specifying the frame types used in BANMAC as listed in Table I. Fig. 5 shows the MAC header (MHR), payload and footer (MFR) details of these frames. The DATA and RSSI-DATA frames are the same as the data frame specified by the standard, except for the three bits of the frame control field.




Fig. 4: IEEE 802.15.4 superframe structure. In the figure τ denotes the Base Superframe Duration.

Packet Type	Frame Control Bits (7-9)	BEM	NBEM
BCN-STD	001	✓	
BCN-NBEM	111		✓
ASSOC-BCN	010	✓	✓
ASSOC-REQ	011	✓	✓
DATA	100	✓	✓
RSSI-DATA	101	✓	✓

TABLE I: Reserved bits' values for the BANMAC frames and applicability to beacon-enabled (BEM) and nonbeacon-enabled (NBEM) modes.

A. Single Network Operation

The network operation begins with the association phase. It is followed by normal operation phase where the (PAN) coordinator collects RSSI time-series and disseminates schedules for collision-free transmissions (Fig. 6). Network set-up spans the association phase and starting of the normal operation phase, where during the latter, RSSI time-series from nodes are collected for the first time. In the association phase, the coordinator periodically broadcasts association beacons, ASSOC-BCN. The periodicity and duration of these broadcasts are determined by two parameters, ABI and ABD –




Fig. 5: BANMAC frames.

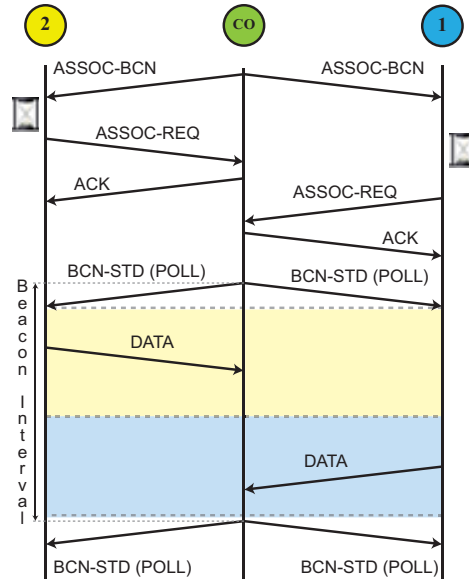


Fig. 6: Exchange of frames: the top part shows the association process.

ABI is the interval between successive broadcasts and ABD is the duration of timeout which starts from the last reception of ASSOC-REQ. The timeout is reset to ABD upon every reception of ASSOC-REQ. Simultaneously, the nodes search for the coordinator by scanning the channels. After a node receives a ASSOC-BCN, it sends the association request in an ASSOC-REQ packet to the coordinator using CSMA/CA. Upon the reception of ASSOC-REQ, the coordinator sends an ACK to the node. After transmitting ASSOC-REQ, a node waits for the reception of the ACK until a user defined timeout. Upon timeout, it sends ASSOC-REQ again.

In the normal operation phase, the nodes access the medium according to the schedule specified by the coordinator. The RSSI based OTW computation procedure requires periodic sampling of RSSI. The periodic transmissions that disseminate the schedule are also used as RSSI probes. These periodic transmissions are sent as the IEEE 802.15.4 standard beacons, BCN-STD, in the beacon-enabled mode and are sent as BCN-NBEM packets in the nonbeacon-enabled mode. For the ease of presentation, we use the common name POLL for these two packet types, which should be disambiguated according to the mode. The uplink transmissions from nodes are done using the same packet format in both modes. In the following, we discuss the issues associated with using the beacon-enabled and nonbeacon-enabled modes.

Beacon-Enabled Mode: As we mentioned earlier, the beacon intervals are limited by the values of the beacon order parameter which takes integer values only. For $BO = 0, 1, 2, 3, 4$, the beacon intervals correspond to approximately 15, 30, 60, 120 and 240 ms, respectively [8]. Fortunately, these values suffice for BANMAC. The step frequency of a walking person is around 1 Hz and that of a running person is 2-3 Hz. Sampling RSSI at 2 to 3 times the Nyquist frequency is

sufficient for the OTW prediction algorithm. Hence, the upper limit on the beacon interval requirements for BANMAC range from 250 ms down to 50 ms. For example, in our evaluations the subjects were walking and for this scenario, $BO = 3$ was sufficient. We note that BO can be set adaptively by first sampling RSSI at high frequency and then adjusting the parameter to the largest value that can satisfy the sampling constraints. The drawbacks of using the beacon-enabled mode are that the beacon intervals cannot be changed dynamically with enough flexibility and that the beacon intervals can take only few values.

Nonbeacon-Enabled Mode: As we mentioned earlier, in the nonbeacon-enabled mode the transmissions are fully asynchronous where the PAN coordinator can communicate with end devices directly *only* using broadcast packets. The coordinator sends periodic broadcasts for RSSI probing and schedule dissemination using BCN-NBEM frames with appropriately formatted payload. In the beacon-enabled mode, the ID of the node designated to collect the RSSI of the beacon is stored. Since in the nonbeacon-enabled mode, such transmissions are broadcast, the bytes 8 and 9 of the data frame are set to the standard broadcast ID, 0xFFFF.

Post-Association Phase: The coordinator starts the process of collecting RSSI time-series for the first time by sending POLLS without payload for a user-configurable time interval parameter (we used 5 seconds for the tests). Nodes send the RSSI data and its sequence identifier aggregated in one or more packets for each received transmission to the coordinator according to the schedule specified by the coordinator. In the event of packet loss, the coordinator simply reschedules this transmission (loss of individual RSSI probes are handled as specified earlier in sec. II). Upon receiving the RSSI time-series from all nodes, the coordinator determines whether a node is on a moving limb, the magnitude of RSSI fluctuations at the node and the center of OTWs. It computes the opportunistic transmission schedules and starts sending schedules in POLLS periodically according to the newly computed step period.

As we mentioned earlier, if the limbs move at-all in normal human movements such as walking or jogging, the left hand and the right leg move in synchrony and so do the other (right hand and left leg) pair. Therefore, to determine the OTWs for all nodes, it is sufficient to collect RSSI time-series from only one node on a mobile limb. The coordinator determines the node responsible for collecting and transmitting RSSI time-series back to the coordinator. Although not implemented in the current version of BANMAC, it is an option to rotate among the nodes the designated node based on energy and load considerations. Since every beacon frame includes the designated node's ID, the frequency of this rotation can be as high as that of POLL transmissions.

The coordinator computes transmission schedule according to user specified admission and prioritization policies. The schedules are specified as the offset and duration in the superframe and are broadcast by the coordinator in the payload

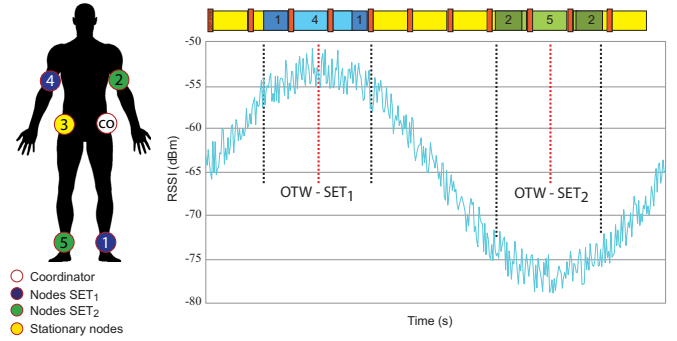


Fig. 7: OTW and transmission schedules. Transmissions from nodes 4 and 5 have higher priority.

of POLL. If a node must first listen for a peer transmission before transmitting to the coordinator, for example to facilitate redundant transmissions for reliability or to route packets when the network topology is not strictly star, the coordinator specifies this in the schedule by setting a forwarding flag and specifying the identifier of the peer node.

B. Co-Located Networks

In parks, nursing homes, public facilities and even in private homes, several BANs can exist together. Furthermore, the BANs in a set of co-located BANs will change dynamically due to movements of people, going from one place to another. Since co-located BANs share the spectrum, managing channel access dynamically to avoid collisions and mutual interference is important. As we discussed in the previous subsection, data transmissions in BANMAC are scheduled by the coordinator. Hence, by assigning a unique channel to each co-located BAN in the same collision domain, we ensure conflict-free medium access globally. In the following, we assume that there exists at-least one free channel for each BAN.

Channel assignment can be done by a global coordinator (GC) in nursing homes and public facilities, and in a fully distributed manner outdoors or in homes and places where it would be unreasonable to assume the existence of GCs. Using GC has the benefit of enforcing certain channel assignment policies suited to the facility in question. A fully distributed coordination does away with the overhead of maintaining a GC, but has the downside of longer convergence time and potentially higher frequency of channel arbitration.

In the following, we outline channel assignment procedures in the presence of a GC as well as in ad-hoc co-location mode. In the presence of a GC, BANMAC seamlessly switches to globally coordinated mode and in the absence of a GC, it automatically switches to decentralized coordinated mode. The channel assignment algorithms presented below are designed to function within the constraints of the IEEE 802.15.4 standard and to meet the resource constraints of the end devices in a BAN. We do not claim these algorithms by themselves to be a novel contribution – our contribution lies in designing ones that are compatible with the IEEE 802.15.4-based BANMAC.

1) *Globally assisted co-location*: We reserve one of the channels for control communications. The GC (global coordinator) transmits standard periodic beacons on the reserved channel in the beacon-enabled mode. In the nonbeacon-enabled mode, it listens for the beacon requests transmitted by BAN coordinators on the reserved channel, to which it answers with its beacon. Then, the usual association process takes place. An important trade-off in using these two modes is that in the beacon-enabled mode, a scan process is a passive listening on a given channel, while in the nonbeacon-enabled mode, the scan process is an active transmission of beacon requests on a given channel.

The GC maintains a database of available channels and current channel assignments. Channel allocations are done using a user defined policy. In our implementation, we use the policy of leasing channels to BANs for a fixed time interval, defined by the parameter *MAX-LEASE*. After the expiration of the lease, the GC puts back this channel in the list of available channels. If the BAN is still in the range of the GC, the local coordinator sends a channel lease renewal request and the process repeats upon each lease expiration. We chose this policy in order to eliminate the need for BANs to transmit liveness status periodically and to free the GC from the burden of scanning channels to determine whether a certain BAN is still present in its range.

The local coordinator of a BAN starts scanning for a GC in any of these three cases: (i) the coordinator starts its BAN, (ii) by carrier sensing BANMAC detects co-channel interference and (iii) performance degradation (i.e., increased packet loss rate). If a BAN coordinator does not detect any GC, it automatically switches to the distributed coordination mode (described in next subsection). Otherwise, it sends a channel allocation request to the GC using the standard IEEE 802.15.4 association mechanism and receives the channel assignment encoded in the 16-bit PAN address. Ideally, the lease times should be short in order to maintain a sufficiently large pool of available channels. On the other hand, a too small lease time will cause large channel allocation overhead.

Since the end devices of a BAN are programmed with the knowledge of the address of the BAN coordinator, they scan all the 16 channels defined by the standard except the control channel to search their pre-configured parent. When they find it, they switch to the channel used by the coordinator and start their operation as in the single-network case described earlier.

2) *Ad-hoc co-location*: In the absence of a GC, BANs converge to a conflict-free channel assignment by running a fully distributed channel assignment algorithm. While the end devices keep the same behavior as described above, every local BAN coordinator starts with a scan process on all the channels, searching for available channels. Out of the available channels, the coordinator picks a random channel and starts a new scan on this channel for a random period. If the channel remains free till the end of this period, it assumes the ownership of the channel for a pre-defined lease time. If at anytime during the (second) scan the channel is found busy, the coordinator marks the channel unavailable, it picks a new random channel




Fig. 8: BANMAC modules.

from the list of available channels and repeats a new scan for a random period. If at the end the BAN coordinator is not able to find any free channel, then it reverts back to the beginning, i.e., it starts a new scan on the control channel searching for a GC and then follows it with repeating the full process of distributed channel assignment. Note that selecting a random channel and following it up with scanning for a random duration is to avoid multiple BANs converging to one channel, which can easily happen if multiple BANs start the process simultaneously, for example, when a sudden interference from e.g., a WLAN occurs.

IV. EVALUATIONS

A. Implementation Remarks

We implemented BANMAC on top of the IEEE 802.15.4 implementation in TinyOS [9]. We used TelosB motes for evaluations. TelosB has the Texas Instruments MSP430 MCU and IEEE 802.15.4 compliant Chipcon CC2420 transceiver. Fig. 8 shows the implementation schema. The end-devices and the local and global coordinators are implemented in TinyOS. The functionalities of BAN coordinator are splitted: due to the limitations of TelosB, the RSSI time-series processing is implemented in C++ and R [10] which we run on a notebook; and the rest in TinyOS. These two modules can be run together on an android platform or on a mote with sufficient computational and storage resources. The TelosB coordinator and notebook are connected by a USB cable.

In our experiments, we measure packet loss rate (PLR), RSSI level, PLR for prioritized transmissions, MAC behavior in co-located networks and the time taken for channel allocation. We compare the performance of IEEE 802.15.4-based BANMAC with the standard IEEE 802.15.4 MAC. In order to keep the experiments tractable without affecting the evaluations of the above mentioned metrics, we used the following configuration:

- Networks operate in beacon-enabled mode.
- The parameters *BO* and *SO* were 3. At this setting, the beacon frequency is suitable for normal human walking (approx. 1 step per second).
- The timeout for association ACK receipt was 864 μ s.
- The OTW prediction algorithm ran every 64 beacon intervals.

B. Single BAN

We evaluated the reliability and differentiated service capabilities of BANMAC. We describe our experimental set-up and summarize the evaluation data in the following.

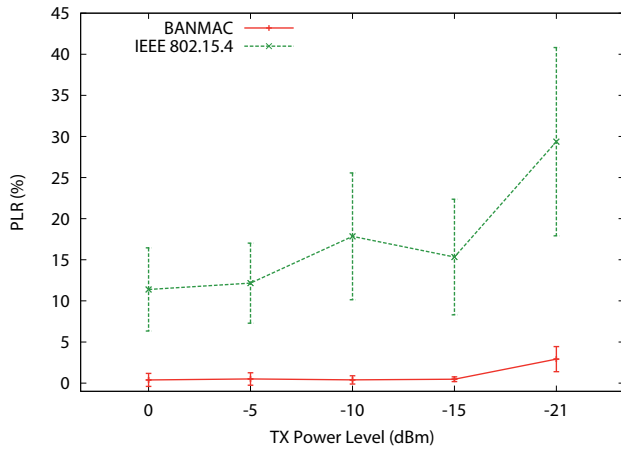


Fig. 9: PLR vs. transmission power.

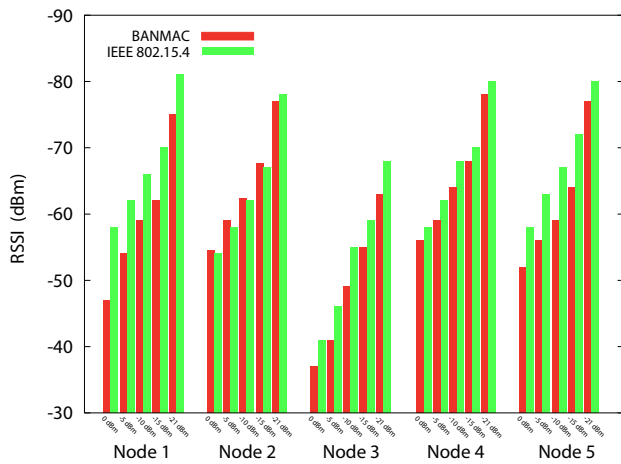


Fig. 10: Average RSSI of data packets.

Reliability: We used a BAN of six nodes. The placement of the nodes over body is shown in Fig. 7 – we placed one node on each arm, each thigh and each ankle of the subject. The nodes on the thighs (kept in trouser pockets) were stationary nodes, one of which is the coordinator. The other four nodes were moving while the subject walked. The coordinator sent beacons every 120 ms with a fixed transmission power of -10 dBm. Each node sent 4 packets/s to the coordinator, which amounts to a load of 20 packets/s at the coordinator. Per-node data point represents approximately 1200 samples. We vary the transmission power of nodes from 0 dBm to -21 dBm and measure the PLR. The data is summarized in Table II and Fig. 9.

Table II shows that BANMAC performs consistently better than the IEEE 802.15.4. The PLR for BANMAC is nearly zero for all configurations except for the case of Tx power set to -21 dBm. In this configuration, the coordinator was losing the RSSI time-series transmissions from nodes. This resulted in errors in OTW predictions. Observe that the BANMAC PLR is still significantly lower for the stationary Node 3 as compared to that of the IEEE 802.15.4 due to conflict-free scheduling of the transmissions. We also measured the RSSI

of transmissions from all nodes to the coordinator. Fig. 10 summarizes the RSSI data. Observe that the height of the bar is inversely correlated to signal strength – taller bar implies lower RSSI. The RSSI of the same node was consistently higher in the case of BANMAC. The differences in the individual RSSI levels are due to differences in the signal attenuation, mainly due to body shadowing. In the light of the stringent reliability requirements of medical applications, the nearly 0% packet loss exhibited by BANMAC affirms the usefulness of our approach. In terms of energy savings, BANMAC significantly outperforms the IEEE 802.15.4. Since 0 dBm is the maximum Tx power, we could not increase the Tx power while evaluating the IEEE 802.15.4 to possibly lower its PLR. The PLR at -21 dBm² for BANMAC is comparable or less than that at 0 dBm for the IEEE 802.15.4. Thus the IEEE 802.15.4 MAC versus BANMAC transmission power ratio for comparable PLR is approximately 125. Fig. 9 shows the PLR for BANMAC and the IEEE 802.15.4 averaged over all tests which reaffirms that BANMAC offers high reliability.

Differentiated Service: Intuitively, since the center of OTW corresponds to the peak of smoothed RSSI fluctuations, the transmissions of a node scheduled close to the center of its OTW should result in lower PLR. We refer to the transmissions scheduled close to the OTW center as higher priority transmissions. To test this hypothesis, we compared the PLR of a node’s transmissions in both higher priority mode and in lower priority mode (basically, we change the scheduling policies at the coordinator to enforce the alternating prioritization at compile time). For this evaluation, we use the same node placement as before, except that we removed the stationary Node 3. We kept transmission power of all nodes and the coordinator fixed at -10 dBm. The reported per-node data point represents approximately 1000 samples.

In table III, we compare the PLR for the transmissions of the same nodes in high and low priority configurations. The data for higher priority transmissions are marked in green color. Fig. 11 shows both the RSSI and PLR statistics. We find that the PLR in the high priority mode is about 50% less than that in the lower priority mode. Furthermore, the RSSI values are almost the same or higher for the high priority transmissions. The RSSI of Nodes 1 and 2 are nearly the same because they were close to the coordinator (all three on the left side). However, we find that the RSSI of nodes 4 and 5 are significantly higher in the case of high priority. In addition to providing the evidence that BANMAC is capable of providing differentiated service, these experiments reaffirm our fundamental premise that scheduling transmissions closer to the OTW centers leads to higher reliability.

To verify the correctness of prioritized scheduling, we connected a digital oscilloscope to the power supply of the nodes. In Fig. 12, we show the traces of the transmissions from *OTW – Set₁*. The two lines in the bottom part of the figure show the schedule in detail in different superframes. The BCN-STDs, which indicate the span of a superframe, are

²-21 dBm corresponds to approximately 0.008 mW.

PLR																				
TX POWER (dBm)	BANMAC										IEEE 802.15.4									
	Node 1		Node 2		Node 3		Node 4		Node 5		Node 1		Node 2		Node 3		Node 4		Node 5	
	Avg. (%)	Std. dev.	Avg. (%)	Std. dev.	Avg. (%)	Std. dev.	Avg. (%)	Std. dev.	Avg. (%)	Std. dev.	Avg. (%)	Std. dev.	Avg. (%)	Std. dev.	Avg. (%)	Std. dev.	Avg. (%)	Std. dev.	Avg. (%)	Std. dev.
-21	2.92	1.81	2.38	2.40	0.69	0.24	3.93	1.21	4.67	0.397	33.25	8.03	34.88	4.50	8.92	6.92	34.81	9.98	34.99	7.94
-15	0.14	0.24	0.70	0.48	0.83	0.42	0.28	0.24	0.42	0.005	13.87	4.56	17.48	4.81	4.10	1.30	18.48	9.84	22.76	5.28
-10	0.14	0.24	0.42	0.42	1.25	0.84	0.14	0.24	0	0	18.58	4.12	21.26	2.28	4.37	2.81	21.90	7.74	23.14	5.24
-5	0.14	0.23	0.57	0.254	1.80	1.26	0	0	0	0	15.98	3.95	14.41	2.81	3.73	0.66	14.15	4.04	12.52	0.91
0	0	0	0.14	0.236	1.80	1.041	0	0	0	0	12.63	3.34	14.36	3.46	2.42	2.22	13.44	6.63	14.10	6.07

TABLE II: PLR statistics

PLR																	
% OTW used	BANMAC – High Priority Nodes: 1 and 2								BANMAC – High Priority Nodes: 4 and 5								
	Node 1		Node 2		Node 4		Node 5		Node 1		Node 2		Node 4		Node 5		
	Avg. (%)	Std. dev.	Avg. (%)	Std. dev.	Avg. (%)	Std. dev.	Avg. (%)	Std. dev.	Avg. (%)	Std. dev.	Avg. (%)	Std. dev.	Avg. (%)	Std. dev.	Avg. (%)	Std. dev.	
80	0.26	0.077	0.18	0.1	0.39	0.22	0.24	0.061	0.46	0.223	0.41	0.230	0.19	0.091	0.11	0.05	

TABLE III: PLR statistics for prioritized transmissions.

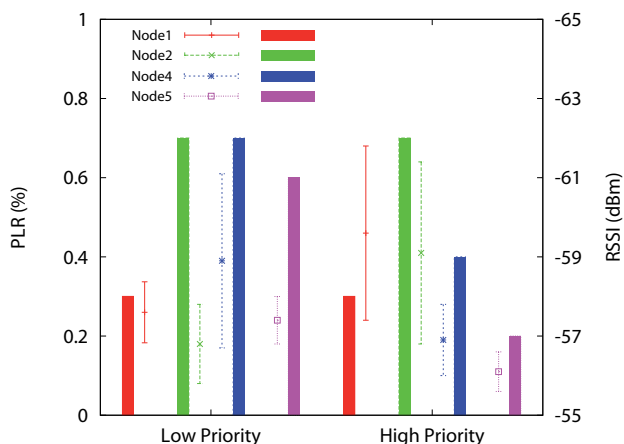


Fig. 11: PLR and RSSI for prioritized transmissions. PLR is shown with lines and RSSI is shown with solid bars.

marked with arrows. The traces on the upper one of these two lines indicate the ON status of ratio and the lower one indicate the OFF status. During the test, we had applied the scheduling policy of higher priority to node with larger ID (Node 4). The trace shows that the transmissions from Node 4, shown in yellow on the upper line, are distributed close to the center of OTW and the transmissions from Node 1 (lower priority) are farther away from it. The start of data transmissions are marked with an asterisk.

C. Co-Located BANs

We evaluated the performance of the centralized and fully distributed channel allocation mechanisms described in Section III-B with the following three test scenarios:

- Test Case 1: 3 BANs, 1 GC (global coordinator) and 2 interferers.
- Test Case 2: 4 BANs, 1 GC and 1 interferer.
- Test Case 3: 5 BANs, no GC and 1 interferer.

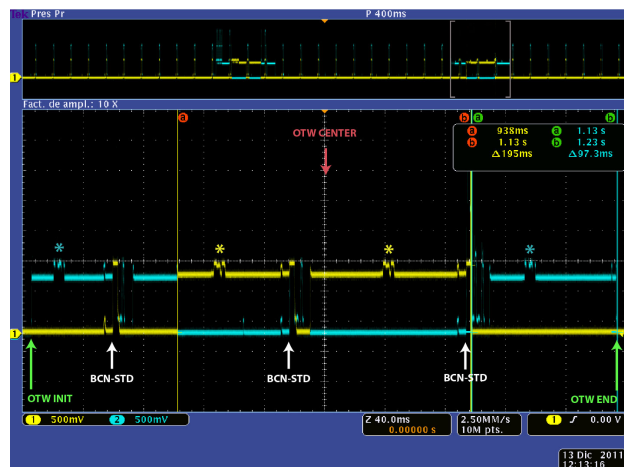


Fig. 12: Prioritized scheduling.

In our experiments the number of co-located BANs was at most 5. Therefore, to test the coordination procedures at full capacity, we restricted the number of available channels from 16 to 6 – the range of available channels was from 21 to 26. The control channel was set to 26. These channels showed no external interference. In test case 3, during the time interval when the interferer was on, the number of available channels was only 4 while the number of BANs was 5. Thus this test case breaks our assumption of the number of available channels being larger than or equal to the number of BANs. The parameter MAX-LEASE was set to 600 beacon counts, which for $BO = 3$, corresponds to approximately one minute.

Each BAN was composed of 3 TelosB motes. We placed one node on each foot and the third TelosB node was the coordinator, placed on the chest and connected through USB to a notebook carried by the subject. The subjects walked in a large room randomly. To switch the evaluation from centralized to distributed coordination, we simply turned off the GC. Each interferer was configured to periodically broad-

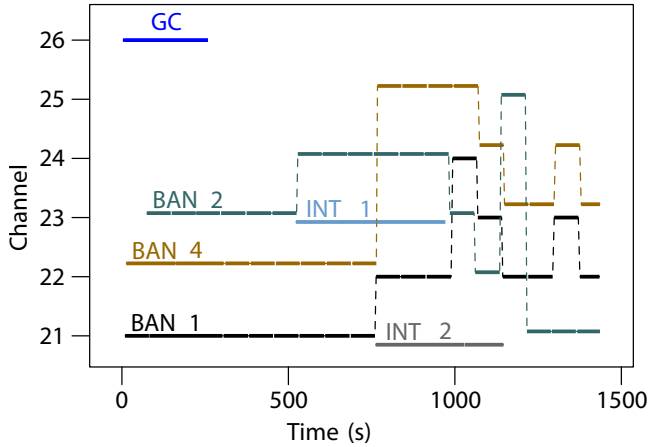


Fig. 13: Channel utilization for Test Case 1. The interferers are configured on channel 21 and 23.

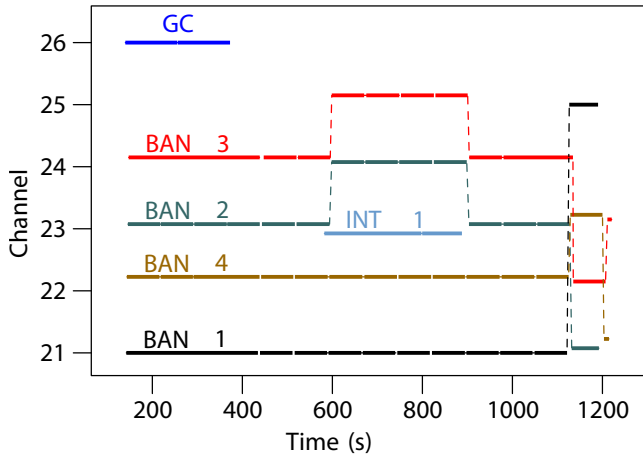


Fig. 14: Channel utilization for Test Case 2. The interferer is configured on channel 23.

cast packets on a fixed pre-defined channel. Specifically, the interferer node was also a TelosB mote that runs similar code as GC and sends beacons with $BO = 2$, but it doesn't accept any association requests. In addition to logging all data sent from the coordinator at the notebooks, we used 6 TelosB nodes that ran a sniffer application, each one on a distinct channel. These nodes grabbed raw IEEE 802.15.4 packets from *all* nodes over the air and sent them to a notebook for logging to which they were connected through a USB hub. We parsed and analyzed the sniffer logs off-line using parsers written in C and MATLAB. The performance metric evaluated in these scenarios is the downtime of each BAN, i.e., the time interval between a channel lease expiration and the granting of new channel, both in the centralized and distributed versions.

Figs. 13, 14 and 15 show the channel utilization for the three test cases. Test cases 1 and 2 were first run with the GC on, followed by the GC switched off and subsequently the interferer(s) turned on. The latter two evaluate the distributed channel allocation mechanism. From the figures, it is evident

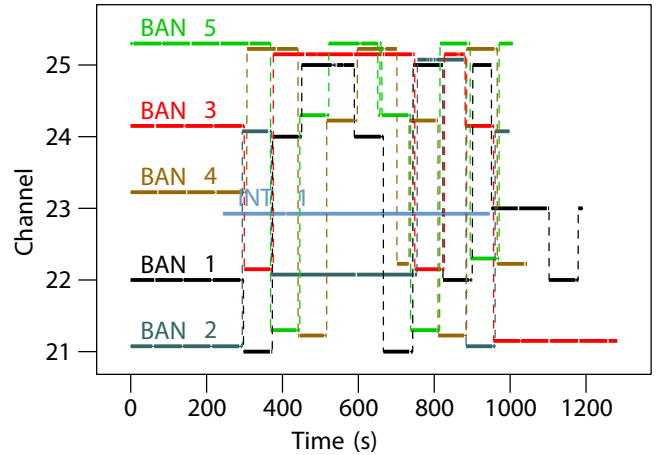


Fig. 15: Channel utilization Test Case 3. The interferer is configured on channel 23.

that the channel allocation is handled well by both mechanisms. The distributed channel assignment reacts quickly when the interferer(s) is(are) turned on and converges to a non-conflicting assignment as long as at-least one channel is available for each BAN.

As we mentioned earlier, Test Case 3 breaks the assumption of at-least one available channel for each BAN when the interferer is turned on. In the beginning, the distributed coordination mechanism assigns conflict-free channels to all 5 BANs. As soon as the interferer starts, the mechanism shows few extra conflicts on channel 25. This is due to a bug in the IEEE 802.15.4 TinyOS implementation, which keeps sending beacons while the node are scanning, even after making a call to the reset MAC function. After the interferer was stopped, Fig. 15 shows that the distributed channel assignment is able to recover to conflict-free channel assignment. Around 1050 s, the batteries of the notebooks started to exhaust and hence the corresponding BANs had to be stopped.

We present the statistics of the channel switching times in Fig. 16. In the case of the centralized coordination mechanism (with the GC on), the channel switching time was approximately 1 s, while in the case of the distributed coordination mechanism (with the GC off) the channel switching time was mostly concentrated in the range of 3 – 4 s. The cumulative distribution function (CDF) plots show that the majority of switching times were below 4 s and the channel switching time is less than 5 s with probability larger than 0.98.

V. RELATED WORK

Maman et al. present BAN channel models using IEEE 802.15.4 MAC [11]. They propose access policies within an IEEE 802.15.4 superframe. The MAC itself is not modified. Another such scheme is presented by Shreshtha et al. in [12], where the authors propose a GTS allocation scheme for BANs with wheelchairs. In [13], the authors present MBStar, a MAC protocol for BANs. MBStar is a TDMA-based MAC and uses encryption for secure transmissions. The goal of this work

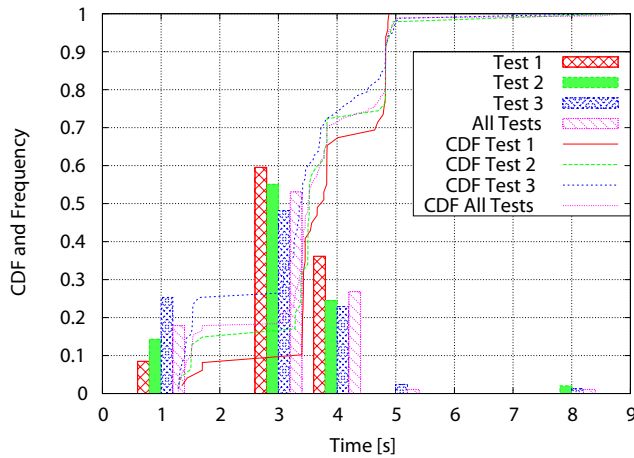


Fig. 16: Channel switching time distribution.

was to achieve high data rate without considering energy consumption. Some of the work presented in the IEEE 802.15.6 Workgroup proceedings is related to this paper [14]. Davenport et al. present a study of link characterization of medical BAN indoors [15]. In [16], Cai et al. derive a two state channel model based on empirical RSSI measurements in BANs, which also match our experimental results. A MAC protocol for BANs is proposed in [17], where throughput maximization is the objective. To the best of our knowledge, our work is the first implementation of an opportunistic MAC that exploits RSSI fluctuations for better reliability.

VI. CONCLUSION

We presented BANMAC, an opportunistic MAC protocol for body area networks. It offers conflict-free medium access by pull-based scheduling. We presented the design and implementation details of IEEE 802.15.4 integrated BANMAC. The protocol supports centralized as well as distributed coordination in the case of multiple co-located BANs. We discussed the trade-offs for the two coordination mechanisms. In the evaluations, we found that scheduling transmissions in the proximity of OTW centers leads to lower packet loss and higher RSSI. This confirms the usefulness of the BANMAC approach as well as enables BANMAC to provide differentiated service. The experimental evaluations show that the packet loss rate of BANMAC is nearly zero unless the transmission power is so low that the nodes can not communicate with the coordinator. An interesting future work is extending BANMAC to provide fair degradation of service for co-located BANs.

ACKNOWLEDGMENTS

We thank Hossein Foutouhi, Athanasios Garyfalos, Abdel Rahman, Manuel Santos and Maryam Wahabi for helping with the experiments. This research was supported by the European Commission under CONET, grant no. FP7-ICT-224053; FEDER funds, grant nos. CSD-2006-00046 and TIN-2009-14475-C04-03; the Spanish MEC and MICINN;

REWIND, grant no. FCOMP-01-0124-FEDER-010050, co-funded by national funds (PT-PIDDAC) through the FCT-MCTES (Portuguese Foundation for Science and Technology) and by FEDER (European Regional Development Fund) through COMPETE; and WSN4QOL, grant no. IAPP-GA-2011-286047.

REFERENCES

- [1] United Nations Department of Economic and Social Affairs, 2011. [Online]. Available: <http://www.un.org/esa/population/publications/worldageing19502050/>
- [2] U.S. Health Resources and Services Administrations, "What is Behind HRSA's Projected Supply, Demand, and Shortage of Registered Nurses?" Sep. 2004. [Online]. Available: <http://bhpr.hrsa.gov/healthworkforce/reports/behindnprojections/index.htm>
- [3] M. A. Hanson, H. C. Powell Jr., A. T. Barth, K. Ringgenberg, B. H. Calhoun, J. H. Aylor, and J. Lach, "Body area sensor networks: Challenges and opportunities," *Computer*, vol. 42, pp. 58–65, 2009.
- [4] K. S. Prabh and J.-H. Hauer, "Opportunistic packet scheduling in body area networks," in *EWSN '11: In Proc. of the 8th European Conference on Wireless Sensor Networks*. Springer LNCS, February 2011, pp. 114–129.
- [5] ZigBee Alliance, 2011. [Online]. Available: <http://www.zigbee.org>
- [6] —, "ZigBee wireless sensor applications for health, wellness and fitness," 2010. [Online]. Available: <http://www.zigbee.org/Standards/ZigBeeHealthCare/Overview.aspx>
- [7] J. D. Troyer, "RSSI based opportune transmission window estimation in unreliable channel conditions," Master's thesis, School of Engineering, Polytechnic Institute of Porto, 2011.
- [8] *IEEE Standard for Information technology Part 15.4: Wireless Medium Access Control (MAC) and Physical Layer (PHY) Specifications for Low Rate Wireless Personal Area Networks*, LAN/MAN Standards Committee of the IEEE Computer Society Std., September 2006.
- [9] J.-H. Hauer, R. Daidone, R. Severino, J. Busch, M. Tiloca, and S. Tennina. (2011, February) An open-source IEEE 802.15.4 MAC implementation for TinyOS 2.1. Poster Session at 8th European Conference on Wireless Sensor Networks. [Online]. Available: <http://www.nes.uni-due.de/ewsn2011>
- [10] R Development Core Team, *R: A Language and Environment for Statistical Computing*, R Foundation for Statistical Computing, Vienna, Austria, 2011, ISBN 3-900051-07-0. [Online]. Available: <http://www.R-project.org>
- [11] M. Maman, F. Dehmas, R. D'Errico, and L. Ouvry, "Evaluating a TDMA MAC for body area networks using a space-time dependent channel model," in *Personal, Indoor and Mobile Radio Communications, 2009 IEEE 20th International Symposium on*, sept. 2009, pp. 2101–2105.
- [12] B. Shrestha, E. Hossain, and S. Camorlinga, "IEEE 802.15.4 MAC with GTS transmission for heterogeneous devices with application to wheelchair body-area sensor networks," *Information Technology in Biomedicine, IEEE Transactions on*, vol. 15, no. 5, pp. 767–777, sept. 2011.
- [13] X. Zhu, S. Han, P.-C. Huang, A. Mok, and D. Chen, "Mbstar: A real-time communication protocol for wireless body area networks," in *Real-Time Systems (ECRTS), 2011 23rd Euromicro Conference on*, july 2011, pp. 57–66.
- [14] "IEEE 802.15 WPAN TG6 body area networks," 2012. [Online]. Available: <http://www.ieee802.org/15/pub/TG6.html>
- [15] D. Davenport, F. Ross, and B. Deb, "Wireless propagation and coexistence of medical body sensor networks for ambulatory patient monitoring," in *Wearable and Implantable Body Sensor Networks, International Workshop on*. Los Alamitos, CA, USA: IEEE Computer Society, 2009, pp. 109–113.
- [16] J. Cai, S. Cheng, and C. Huang, "MAC channel model for WBAN," IEEE P802.15, Tech. Rep. 15-09-0562-00-0006, July 2009.
- [17] S. M. Yoo, C. J. Chen, and P. H. Chou, "Low-complexity, high-throughput multiple-access wireless protocol for body sensor networks," in *Wearable and Implantable Body Sensor Networks, International Workshop on*. Los Alamitos, CA, USA: IEEE Computer Society, 2009, pp. 109–113.