Trade-off Analysis of a MAC Protocol for Wireless e-Emergency Systems

Óscar Gama[1], Paulo Carvalho[1], J.A. Afonso[2], and P.M. Mendes[2]

[1] Department of Informatics,
University of Minho, Braga, Portugal
{osg,pmc}@di.uminho.pt
[2] Department of Industrial Electronics,
University of Minho, Guimarães, Portugal
{jose.afonso,paulo.mendes}@dei.uminho.pt

Abstract. Wireless sensor networks are envisioned to be deployed in healthcare. Since emergency and intensive care applications need to assure reliable and timely data delivery, they have increased demands for quality of service, including at the MAC layer. Amongst MAC protocols available for WSNs, the Low Power Real Time (LPRT) presents suitable characteristics to be deployed in emergency platforms due to its rational bandwidth allocation, low energy consumption, and bounded latency. Yet, this protocol may present a significant packet loss ratio in a wireless channel with bit error ratio. In order to define a MAC protocol more robust to bit error conditions and able to fulfill the required quality of service, solutions based on short size beacons and multiple retransmissions are proposed and tested. The results showed that such strategies led to meaningful improvements regarding packet loss ratio, without compromising significantly the energy consumption.

Keywords: e-Emergency, Quality of Service, Wireless Sensor Networks.

1 Introduction

An e-health monitoring system commonly grounds on a group of sensors attached non-invasively to a patient in order to monitor some physiological parameters. In case of emergency clinical scenarios, a healthcare network should provide quality of service (QoS) facilities since these clearly demand for high reliability, guaranteed bandwidth and short delays [1]. Therefore, communication protocol layers need to assure a reliable and timely data delivery.

Many Medium Access Control (MAC) protocols have been developed for wireless networks using contention or multiplexing-based algorithms. Traditional contention-based protocols assume traffic is distributed stochastically. As traffic in a wireless sensor network (WSN) tends to be highly correlated and dominantly regular, conventional Carrier Sense Multiple Access (CSMA) protocols are not advised for WSNs [2]. S-MAC [3] and WiseMAC [4] are typical examples of CSMA-based protocols designed for low duty-cycle WSNs, developed to help saving energy in applications whose nodes remain idle for long time until an event is detected (e.g. surveillance).

S. Hailes, S. Sicari, and G. Roussos (Eds.): S-Cube 2009, LNICST 24, pp. 222–235, 2009.

These protocols are convenient for WSNs having usually low traffic loads. Consequently, these MAC protocols are inadequate for networks requiring high throughput and low latency.

Amongst multiplexing-based protocols, Time Division Multiple Access (TDMA) is a commonly used technique. Time is divided into slots, which are used by motes (wireless sensor nodes) to transmit data without the need to contend for the medium. If a base-station (BS) is available to keep the WSN scheduling and synchronization, then TDMA-based MAC protocols are usually a preferable choice to satisfy the QoS requirements of e-emergency systems, as QoS is easier assured in a collision-free environment than in a contention-prone medium. Within TDMA-based MAC protocols, LPRT [5] is a convenient choice to provide efficient bandwidth allocation, low energy consumption, and bounded latency, as required by e-emergency wireless networks. However, it may present a significant packet loss ratio in a wireless channel affected by errors. In order to improve its robustness to bit errors, new solutions based on short size beacons and multiple retransmissions are proposed and verified in a WSN simulator.

The remaining of this paper is organized as follows; the related work is discussed in Section 2; proposals to improve LPRT protocol are presented in Section 3; the simulation testbed setup is described in Section 4; the results are discussed in Section 5; and, finally, the conclusions are presented in Section 6.

2 Related Work

TDMA-based protocols available for WSNs include IEEE 802.15.4/GTS [7], LMAC [8], TRAMA [9], and LPRT.

The IEEE 802.15.4 standard specifies the physical layer and the MAC sublayer, allowing the optional use of a TDMA-based superframe structure. However, the low granularity of the time slots leads to poor bandwidth efficiency, making it unsuitable to e-emergency scenarios [1].

LMAC allows a WSN to self-organize in terms of slot assignment and synchronization, without requiring a central manager, through a distributed algorithm running in every mote. Nodes wake up at the start of each slot to stay synchronized and to listen to a message. If the message is not addressed to the node, it will sleep until the next slot. Since each node has only one slot assigned and one transmission per superframe (32 slots) allowed, LMAC is inadequate to e-emergency systems.

TRAMA uses a distributed election scheme based on information about the traffic at each node to decide which node can transmit at a particular slot. It avoids the assignment of slots to nodes without traffic to send, and also allows nodes to decide when they can sleep. It is well suited for applications that require high-delivery guarantee and energy efficiency, without being delay sensitive. The latter characteristic impairs the support of real-time applications.

LPRT is a simple, beacon-based protocol that uses dynamic and efficiently the available bandwidth. Its highly-grained superframe (Figure 1) starts with the transmission of a beacon frame (B) by the BS, followed by the Contention Access Period (CAP), also called Contention Period (CP). The CAP may be used for the (dis)association or configuration of a body sensor network (BSN). The Contention

Free Period (CFP) follows the CAP. The CFP is composed by the Normal Transmission Period (NTP) and the Retransmission Period (RP). The NTP is used for motes to transmit new data. Lost data are retransmitted in the RP. Data packets are sent in contiguous slots of the CFP. The slots attribution in the CFP is announced through a list of allocation fields carried in the payload of every beacon. Each allocation field contains the association identification and the initial transmission slot (ITS) for every mote in the WSN. Data frames transmitted to the BS during the NTP are acknowledged by the ACK bitmap present in the beacon of the next superframe.

The beacon size may become relatively large, since it is directly dependent on the number of motes associated with the BS. A mote transmits data in the superframe only if the corresponding beacon carrying the list of allocation fields is received, and a single retransmission procedure is used in case of transmission failure. These characteristics may lead to a significant packet loss ratio if communications occur in a wireless channel with an appreciable bit error ratio (BER). In order to improve its robustness against bit errors, a solution based on short size beacons is proposed.

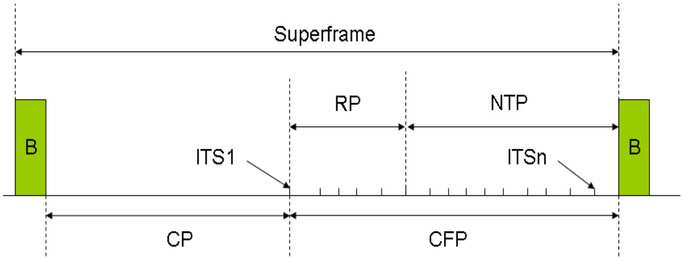

Fig. 1. Superframe structure in the LPRT protocol [5]

3 Proposals to Improve LPRT

In order to define a more robust MAC protocol to channel errors, solutions based on short size beacons and multiple retransmissions were added to LPRT. For this purpose, the RP is placed between the CAP and the NTP. RP is defined after the CAP so that a packet is retransmitted away from an eventual burst error condition responsible for the transmission failure occurred during the last NTP. RP is not placed after the NTP to avoid a mote transmitting in the NTP with a variable packet size.

Short size beacons. To assure a good performance of the e-emergency WSN, the percentage of lost beacons should be very low. A strategy to accomplish this goal is to send beacons with a convenient transmission power, since the BER of the channel decreases as SNR increases. In addition, the beacon frame length should be as small as possible. So, whenever possible the beacon payload contains only the ACK bitmap to acknowledge the frames correctly received during the NTP of the last superframe, and the CAP size of the current superframe (assuming the start slot of the CAP is known). In this case, motes must run an algorithm to compute which slots should be used to (re)transmit data without interfering each other, in accordance with a predefined order schema. Using this strategy, the energy consumption in each BSN improves too, since smaller size beacons are received by the motes.

According to the received ACK bitmap, each mote must calculate the corresponding superframe slots to transmit its data. If a mote does not receive a beacon or a short sequence of beacons, it may continue to send its new data in the NTP, since a mote clock drift in the order of microseconds should permit the WSN to continue synchronized during a few consecutive beacon intervals. However, the mote cannot retransmit any data in the RP because the ACK bitmap is not available and so it does not know how the RP slots are being allocated to the motes.

In order to save energy, it might be tempting to retransmit the lost data aggregated to the new data sent in the NTP, instead of retransmitting it in the RP. However such strategy should be avoided because the number of slots allocated in the NTP to each mote becomes variable. Consequently, if a mote does not receive the beacon, it has no way to know *a priori* which slots to use for transmission. Indeed, if a mote does not receive the ACK bitmap, it does not know which motes are going to transmit aggregated data in the current superframe, making impossible to compute the new slot allocation schema. This situation does not occur if the NTP is used only to transmit new data packets. Hence, aggregated retransmissions in the NTP are not recommended, except in cases where aggregation does not imply taking more slots, such as packets carrying temperature data.

Besides the ACK bitmap and the CAP size, a beacon may need to send reconfiguration instructions if a new clinical situation is detected in some BSN. For instance, higher monitoring activity and lower delay transmission of the vital signals might be required when a patient´s clinical situation changes from non-critical to critical. In this case, the BS should inform all motes about the new situation and eventually reconfigure the WSN. To perform this action all motes must follow a reconfiguration scheme, such as the algorithm proposed in [6].

Retransmission in the RP. The retransmission order in the RP depends on the ACK bitmap received from the BS. Using an increasing slot sequence, firstly data of all motes having the highest-priority and the bit false in the ACK bitmap are retransmitted successively. Then, data of all motes having the second higher-priority and the bit false in the ACK bitmap are retransmitted successively, and so on. If slots are not enough to permit all required retransmissions, then schedule truncation is done in order to guarantee that no retransmission occurs in the NTP. It should be noted that retransmissions may not be the appropriate error recovery mechanism if losses are due to fading.

Retransmission in the CAP. If a mote does not receive the beacon then it fails to receive the ACK bitmap, and therefore it does not know if the BS correctly received the packet sent in the NTP of the last superframe. Retransmitting properly that packet in the RP is impossible because such mote cannot compute the slot allocation schema of RP. A solution to overcome this problem is to retransmit it during the CAP using the slotted CSMA/CA [7]. This procedure should improve the packet loss ratio at expense of some energy consumption and CAP slots waste, since the mote may be transmitting a duplicated packet. Indeed, a packet already received by the BS may be retransmitted again in the CAP of the next superframe if the mote does not receive the beacon.

4 Simulation Testbed

To test the solutions proposed in the previous section and to compare them with LPRT, the Castalia simulator [10] was used. For that purpose, a case-study and distinct MAC protocol operation modes were programmed in Castalia.

4.1 Castalia Simulator

Castalia is a discrete event-driven simulator designed specifically for wireless sensor networks. It uses the communication model proposed by Zuniga *et. al.* [11]. In a wireless channel the BER is variable because errors may occur in long bursts due to fading or shadowing. Castalia does not support currently these time-varying effects.

Castalia provides a tuneable CSMA-based MAC protocol in order to model several link protocols. However this protocol was built for broadcast communications, not unicast, and so it does not currently support acknowledgements. For this reason, acknowledgements are performed at the application level.

In order to compute the total energy consumed by a mote, Castalia takes into account not only the consumed power while the radio is in listening, sleeping, transmitting or receiving state, but also the power consumed by state switching. It also accounts the energy consumed per sample by the sensing devices present at each node. CPU energy consumption and memory access cost are not accounted in the current version of Castalia.

4.2 Case-Study

The testbed is based on a hospital room containing six beds with one patient per bed. Each patient is monitored by a body sensor network, and a BS collects and analyses the vital signals of all patients. The signals being monitored are temperature (T), oximetry (OXI), arterial pressure (ART), respiration rate (RR) and electrocardiography (ECG). Each signal is collected and transmitted by a dedicated mote. The MAC protocol deployed in the network must guarantee a maximum latency below 500 ms, low packet loss ratio, 10.5 kbps of bandwidth to each BSN and low energy consumption [1].

4.3 Operation Modes

In order to compare different MAC strategies, five operation modes were defined in the simulation testbed: 0, 1, 2, 3, and 4. In all modes, frames transmitted at the NTP are acknowledged through the ACK bitmap sent in the payload of the next beacon. Also, a mote may transmit data in the NTP even if a short size beacon is not received.

In operation mode 0, there are no retransmissions of lost frames. This mode is used as the worst-case, as no error control mechanism is present in the system.

In operation mode 1, lost packets are retransmitted once at the RP and these retransmitted packets are not acknowledged. This mode differs of LPRT, because in this protocol a mote only transmits data in the NTP if a relatively large size beacon is received. Hence, if the channel BER is not null, the packet loss ratio in a WSN operating at mode 1 is lower than using LPRT.

In operation mode 2, packets with a payload size above a fixed threshold may be retransmitted during RP at most twice. The threshold was set to 40 bytes, so ECG and ART motes have two chances to retransmit the lost data. The second retransmission occurs only if the frame is not correctly received by the BS during the first retransmission. Hence, the first retransmission must be acknowledged, although the second one does not need to be. Packets with a payload size below the threshold may be retransmitted once in the RP and they are not acknowledged.

In operation mode 3, packets with a payload size above a fixed threshold may be retransmitted at most three times in the RP.

Operation mode 4 is similar to operation mode 2, but if a mote does not receive a beacon, it tries to send in the CAP the packet transmitted in the NTP of the last superframe. A minimum CAP size is available in all superframes.

4.4 Testbed Setup

The relevant setup and parameterization steps followed in the simulation testbed are presented next.

Wireless Channel. Since the room considered in our case-study is relatively small (10mx10m), the BER given by the simulator was always null, meaning that all motes worked in a good connected region. In order to test the robustness of the diverse protocols to different BERs, an additional random error generator was introduced in the wireless channel module of Castalia to force the degradation of the bit error probability.

Radio Model and Operating States. In the present protocol implementation, a mote radio is normally in the sleeping state. It switches automatically to the transmission state whenever it has a packet to send and enters in listening mode five slots before the ending of the superframe to receive the next beacon. The BS never enters in the sleeping state because conserving energy is not a concern for this device.

The CC2420 radio model, used by the popular TelosB motes, and a transmission power level of -10 dBm were chosen.

Beacon Interval and Superframe Time-Slots. The number of slots in the superframe should be as high as possible to tune accurately the time division allocated to each mote and so minimizing the bandwidth waste, but not too high so that the slot duration is beyond the mote timer resolution. Since this is typically in order of microseconds, it was chosen 512 slots per superframe, which means that the duration of each slot is around 0.43ms, considering a beacon interval of 0.220s. This period was chosen for ECG packets to be sent with the payload almost fully-loaded to save energy, and it must not be above 0.250s to respect the maximum delay of 0.500s imposed to all signals, including the retransmitted frames [1].

Physical, MAC, and ACK Frames. Since many commercial motes use ZigBee, a physical layer frame having a total maximum size of 133 bytes (B) was assumed, the same length as specified in IEEE 802.15.4. Its physical header of 6B was also adopted: preamble sequence (4B), start frame delimiter (1B), and frame length (1B).

A MAC header of 6B was tailored to this case-study, containing the fields: frame control (1B), sequence number (1B), destination address (1B), source address (1B), WSN identification (1B), and frame check sequence (2B).

The ACK frame has a MAC header of 6B and payload size null.

Data Packet Payload. ECG motes sample the physical signal at 250 Hz, ART motes at 120 Hz, OXI motes at 60 Hz, RR motes at 20 Hz, and T motes at 2 Hz [1]. The samples of every mote have a resolution of 16 bits. Consequently, packets transmitted from ECG, ART, OXI, RR, and T motes present a payload of 110B, 54B, 28B, 10B, and 2B, respectively.

Transmission in the NTP. The number of slots N_s each mote occupies in the NTP of the superframe to transmit data is:

$$N_s = \mathrm{ceil}(\, t_{tx} * S / t_{SD}\,) \tag{1}$$

where t_{tx} is the transmission duration (in seconds), S is the total number of slots in the superframe, and t_{SD} is the superframe duration (in seconds). The ceiling function ceil(x) returns the integer part of the argument rounded up. For a packet with a physical header P_h, a MAC header M_h, a MAC payload length M_p bytes, and a transmission rate R bps:

$$t_{tx} = (P_h + M_h + M_p) * 8 / R \tag{2}$$

Considering a null overhead for the layers above the MAC layer,

$$M_p = t_{SD} * F * r / 8 \tag{3}$$

where F is the sample rate (samples/s) of a specific mote type, and r is the sampling resolution in bits. As S=512 and R=250 kbps, each ECG, ART, OXI, RR, T frame takes respectively 10, 5, 3, 2, and 1 slots in the superframe.

When allocating each set of slots in NTP to the motes, two additional slots are included for safeguarding purposes. The slots in the superframe are occupied in the following order:

Beacon | CAP | RP | T(5-0), RR(5-0), OXI(5-0), ART(5-0), ECG(5-0) |

––––––––––––––– NTP –––––––––––––––

where ECG(5-0)=ECG(5,4,3,2,1,0) represents the following transmission sequence in NTP: after ECG mote of BSN 5 (ECG5) transmitting its packet, then ECG4, ECG3, ECG2, ECG1, and ECG0 transmit successively their data. The same criterion is applied to the remaining types of motes.

Retransmission in the RP. The retransmission order in the RP depends on the ACK bitmap received from the BS. Using an increasing slot sequence, firstly the data of ECG motes having the bit false in the ACK bitmap are retransmitted successively. Then, if the respective ACK bits are false, ART motes retransmit their data, followed by the OXI motes, the RR motes, and finally the T motes. The number of slots allocated in RP to the motes includes two slots for safeguarding purposes, plus two slots if ACK receiving is required.

In case of no more slots available to be allocated for retransmission, the less important vital signals should not be retransmitted. Since body temperature changes slowly along time, this signal is the first to be discarded in such situation.

One, two or three consecutive set of slots may be used for a single retransmission, according to the operation mode used. In mode 2, the second set of slots is used for retransmission if the packet is not correctly received by the BS during the first retransmission. Accordingly, the first retransmission must be acknowledged. If a packet is sent with success during the first retransmission, then the slots reserved for the second retransmission are unused, resulting in bandwidth waste. The second retransmission is not acknowledged. The same reasoning is applied for mode 3.

Retransmission in the CAP. In operation mode 4, a minimum CAP size of fifty slots is defined for every superframe so that motes have a chance to retransmit data in case of not receiving a beacon. To free some slots, T motes are not allowed to retransmit data in the CAP. The remaining operation modes have no minimum CAP size.

5 Simulation Results

In order to study the improvement obtained in LPRT when short size beacons and multiple retransmissions are deployed in this protocol, several simulation tests were carried out for different operation modes. Packet loss ratio, energy consumption, CAP availability, and scalability were the parameters under test. The performance of a WSN using beacons with different sizes was tested as well.

For each test run, a simulation time of one hour was defined. Extending the simulation time would not affect significantly the results. During that time period, the BS sends around 16363 beacons, carrying in the payload only the ACK bitmap (4B). It is assumed that BER is equal in both communication directions.

In some graphics the probability P of a fully-loaded packet to be received at destination is used instead of the BER. Both parameters are related by:

$$\text{BER} = 1 - P^{1/(8*\text{MPS})} \tag{4}$$

where MPS is the maximum physical frame size. Considering MPS=133B, the BER changes between 0 and around $6.5*10^{-4}$ as P decreases from 1 to 0.5 along the simulation runs. Typical values for BER in a good real wireless channel are of the order 10^{-5} and in a bad channel may be less than 10^{-3}.

5.1 Performance of a WSN with Short Size Beacons

The beacon size should be as small as possible to improve the performance of the WSN. In order to evaluate the impact of the beacon size in a WSN, the loss ratio of ECG packets for the beacon payloads of 4B (the minimum size to contain a bitmap for 30 motes, as required by the case-study), 36B, 68B, and 100B, were tested in the testbed using operation mode 2. The results are presented in Figure 2. It is observed that as the beacon payload becomes larger, the beacon loss probability increases, and so the packet loss ratio.

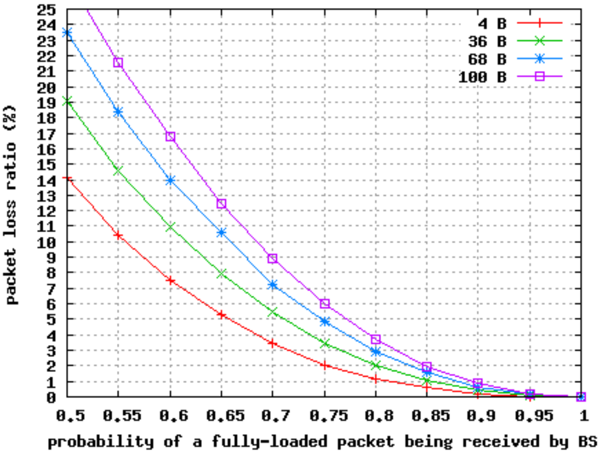

Fig. 2. ECG packet loss ratio for several beacon payload sizes (in steps of 32B)

Moreover, considering that samples performed by sensing devices consume no energy, simulations show that, for operation mode 0 and P=1, ECG motes save around 14% of energy when the BS sends short size (4B) beacons instead of sending in the beacon payload all information (68B) specifying concretely which slots the sensors should take in the CFP.

These results confirm that the performance of a WSN improves significantly in terms of energy saving and packet loss ratio if short size beacons are used.

5.2 Packet Loss Ratio

LPRT may present significant packet loss ratio in wireless channels with bit errors. In order to study how the proposals presented in Section 3 may help to improve this drawback, the Packet Loss Ratio (PLR=100-Packet Delivery Ratio) for the traffic generated by all types of motes were obtained for every operation modes.

Figure 3 presents the packet loss ratio curves for the traffic produced by both the ECG motes and the T motes. The information presented for ECG traffic is calculated by taking into account all packets generated by all ECG motes in the WSN. This criterion holds true also for the other mote types, such as ART motes whose loss ratio curves are presented in Figure 4, and also for OXI and RR motes (Figure 5). By doing so, one may evaluate how the traffic produced by the diverse mote types is affected by operating in a given mode. Yet in a real situation, attention should be paid to the traffic coming from each individual BSN since the traffic treatment may be different for each BSN, according to the degree of the patients' emergency state. However, for easiness of dealing with the whole information produced by the simulator, a global analysis for each type of traffic generated in the WSN is done. This procedure will not affect the main conclusions obtained from the simulation results.

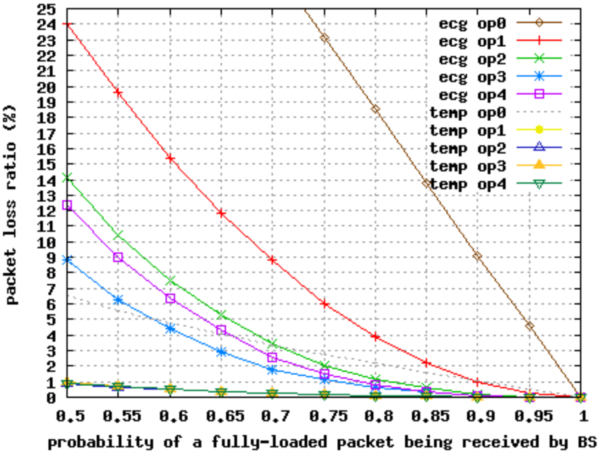

Fig. 3. Packet loss ratio for ECG and T traffic

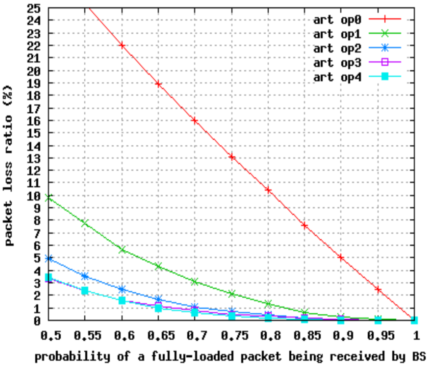

Fig. 4. Packet loss ratio for ART traffic

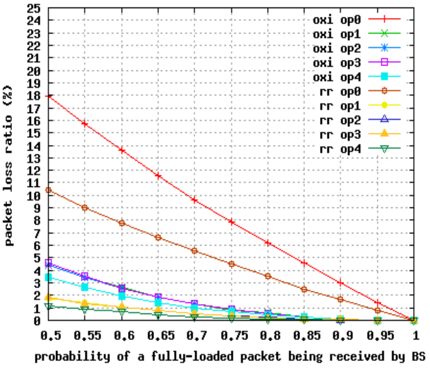

Fig. 5. Packet loss ratio for OXI and RR traffic

Comparatively to mode 1, Figures 3 and 4 illustrate clearly the improvement achieved when the WSN operates in modes 2, 3, or 4. These operation modes present an ECG packet loss below 1% when probability P > 0.8. Such improvement is due to the multiple retransmission process occurred in the superframe. The ECG traffic presents a lower loss ratio in a WSN operating in mode 3 than in mode 4. However for ART traffic it is almost irrelevant to work in mode 3 or 4. This occurs because the probability of losing an ART packet is lower than the probability of losing an ECG packet, as the packet size is smaller in the former.

Since OXI, RR, and T motes may only retransmit once in RCFP, no significant difference is detected in such traffic by operating in modes 1, 2, or 3. However an improvement is noticed at OXI and RR traffic by operating at mode 4. T traffic does not have such improvement because it cannot make use of the CAP.

5.3 Energy Consumption

Since WSNs are operated by low-capacity batteries, energy preservation is important to guarantee that the e-emergency WSN operates for a long time. In order to evaluate the impact of the diverse operation modes on the energy consumption, the lifetime of the motes was studied.

To better evince the impact of each operation mode on the energy cost of the WSN, the consumption due to samplings performed by the sensing devices is ignored initially. For every operation mode, Figure 6 presents the average lifetime per initial energy of the battery for each mote type relative to the simulations run with a probability P=0.75. To know how many minutes a type of mote would live in such conditions, the values in the y-axis must be multiplied by the available energy in the battery (in Joules) when the mote started working. The available energy E in a battery with n cells, having each cell V Volts and a capacity of C mAh, is: E=n*V*C*3.6 J. For instance, one 3V, 1000 mAh lithium coin battery (CR2477) contains 10800 J of initial energy. So, according to Figure 6, an ECG mote operating in mode 4 and powered by that battery would live for 3.68*10800 min = 662.4 hours.

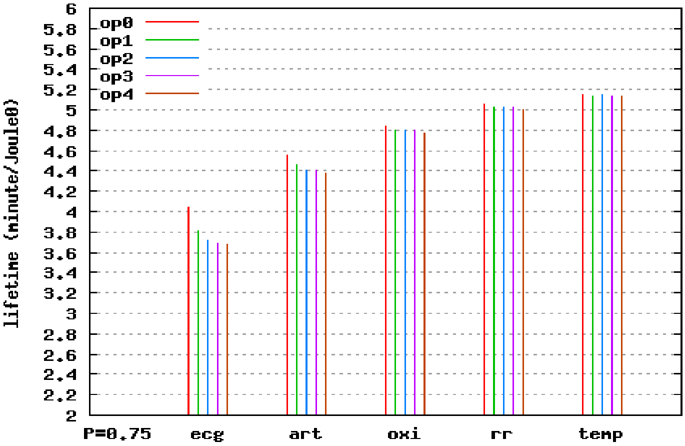

Fig. 6. Lifetime ignoring the sensing consume

For every mote type, the energy consumption does not change significantly with probability P > 0.75 and operation mode 1, 2, 3, or 4.

Figure 7 presents the results when each sampling performed by a sensing device consumes 0.01 mJ of energy. As shown, the effect of the sampling rates on the lifetime of the motes is notorious. As ECG motes have both the highest sampling rate and the largest transmitted packet size, the energy drainage is faster than in the other mote types. However, the lifetime differences observed in each mote type for the diverse operating modes are almost negligible.

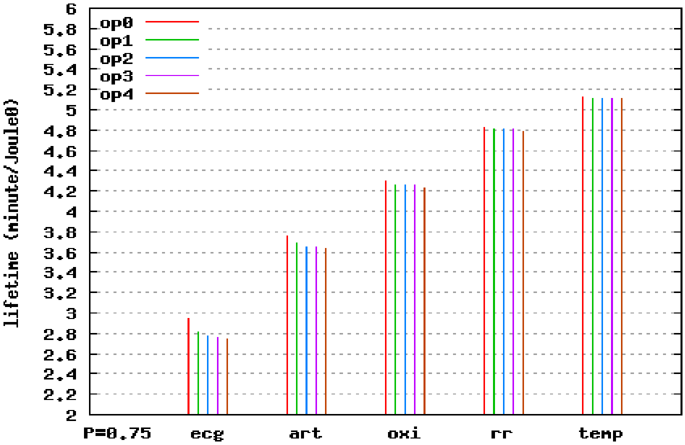

Fig. 7. Lifetime for a sensing consume of 0.01mJ

5.4 CAP Availability

The association or disassociation of a BSN in the WSN is done during the CAP. To see how long is the CAP available for these operations, Figure 8 shows the percentage of superframes containing a CAP with a given size using P=0.75 and operation mode 3, the most limited mode regarding CAP size. In the y-axis, S_1 is the number of superframes with n free slots. For example, a CAP with 255 slots is available in 6.45% of all superframes. It is also shown that information considering groups of slot intervals (bar graph). The values are calculated performing the discrete integral along the specified slot interval. In this case, the values in y-axis must be multiplied by ten, and S_{intv}

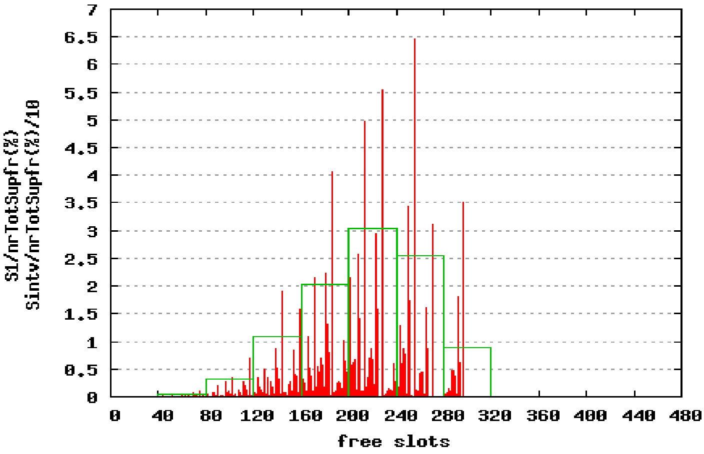

Fig. 8. Free slots for operation mode 3, P=0.75

is the number of superframes with free slots in each interval. For instance, 30.5% of all superframes have a CAP size comprised between 200 and 239 slots.

It is observed that 99.3% of all superframes contain a CAP size above 80 slots, which are enough to perform (dis)association operations. Since CAP size is variable, the beacon should carry the superframe CAP size to inform the motes wishing to associate to the WSN about the available contention slots. It is assumed that the start slot of the CAP is fixed and known previously.

5.5 Scalability

The WSN must be scalable in order to admit additional BSNs. In a LPRT-based system, scalability may be evaluated from the CAP available beyond the minimum CAP size. Indeed, if all superframes have a CAP size bigger than the minimum CAP size, it means that RP and NTP do not take all slots, which eventually may be used by an additional BSN. Figure 9 presents the percentage of superframes with a CAP size equal to the minimum value for several operation modes. This minimum value is fifty slots for operation mode 4, and null for the remaining modes. It is observed that a WSN operating in modes 1, 2, 3, and 4 may admit respectively 5, 3, 1, and 2 additional BSNs without losing significant performance.

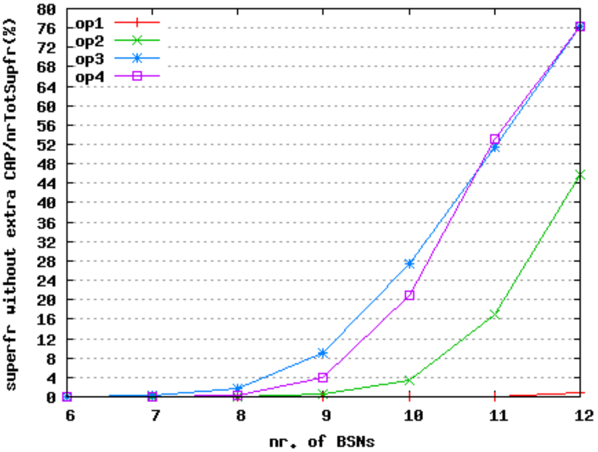

Fig. 9. Evaluation of scalability for P=0.75

5.6 Latency and Goodput

Since all packets are delivered to the BS with a delay below 500ms, the graphs reporting this parameter are not presented. In this slotted framework, goodput is correlated with the packet loss ratio, therefore goodput graphics are not shown too.

6 Conclusions

The deployment of LPRT in e-health systems leads to low power consumption, controlled latency, and throughput efficiency. However, as simulations have shown, LPRT is unreliable if the wireless channel is affected by appreciable bit errors. In order to define a MAC protocol more robust than LPRT, different approaches based on short size beacons and multiple retransmissions have been proposed and tested. The results have shown that such approaches lead to meaningful improvements regarding packet loss ratio, without compromising significantly the energy consumption.

Despite of presenting the best performance regarding packet loss ratio, a network operating in mode 3 faces scalability problems, since significant waste of time slots occurs in the RP. The results have revealed that operation mode 4 offers a better scalability and a good compromise in terms of packet loss. Therefore, we believe that LPRT operating in this mode will enhance QoS in e-health networks. Currently, we are implementing this operation mode in a real testbed for experimental analysis.

Acknowledgments. Óscar Gama is supported by FCT (SFRH/BD/34621/2007), Portugal.

References

1. Gama, O., Carvalho, P., Afonso, J.A., Mendes, P.M.: Quality of service in wireless e-emergency: main issues and a case-study. In: Proc. of 3rd UCAmI, Salamanca, Spain (October 2008)
2. Akyildiz, J.F., Su, W., Sankarasubramaniam, Y., Cayirci, E.: A survey on sensor networks. IEEE Communications Magazine, 102–114 (August 2002)
3. Ye, I., Heidemann, J., Estrin, D.: MAC with coordinated adaptive sleeping for wireless sensor networks. IEEE/ACM Trans. Networks 12(3) (June 2004)
4. El-Hoiydi, A., Decotignie, J.D.: WiseMAC, An ultra low power MAC protocol for the WiseNET wireless sensor network. In: Proc. of 1st ACM SenSys Conf., USA (November 2003)
5. Afonso, J.A., Rocha, L.A., Silva, H.R., Correia, J.H.: MAC protocol for low-power real-time wireless sensing and actuation. In: 13th IEEE International Conference on Electronics, Circuits and Systems, Nice (December 2006)
6. Gama, O., Carvalho, P., Afonso, J.A., Mendes, P.M.: An Improved MAC Protocol with a Reconfiguration Scheme for Wireless e-Health Systems Requiring Quality of Service. In: 1st Wireless Vitae 2009, Aalborg, Denmark (May 2009)
7. IEEE Std 802.15.4-2003, Wireless Medium Access Control (MAC) and Physical Layer (PHY) Specifications for Low-Rate Wireless Personal Area Networks (October 2003)
8. van Hoesel, L.F.W., Havinga, P.J.M.: A lightweight medium access protocol (LMAC) for Wireless Sensor Networks. In: Proc. 3rd Information Processing in Sensor Networks, Berkeley (April 2004)
9. Rajendran, V., Obraczka, K., Garcia-Luna-Aceves, J.J.: Energy-Efficient, Collision-Free MAC for Wireless Sensor Networks. In: Proc. ACM SenSys 2003, LosAngeles, USA (November 2003)
10. Castalia Simulator, `http://castalia.npc.nicta.com.au`
11. Zuniga, M., Krishnamachari, B.: Analyzing the transitional region in low power wireless links. In: 1st IEEE Annual Conference on Sensor and Ad Hoc Communications and Networks (October 2004)