An Ultra-Low Power MAC Protocol for In-body Medical Implant Networks

Ashutosh Ghildiyal[1], Balwant Godara[2], and Amara Amara[2]

[1] Parc d'affaires noveos, 4 avenue Réaumur, 92140 Clamart, France
[2] Institut Supérieur d'Electronique de Paris,
28 Notre Dame des Champs,76005 Paris, France
{ashutosh.ghildiyal}@sorin.com
{bgodara,amara.amara}@isep.fr

Abstract. We present an ultra low power MAC designed for battery-operated subcutaneous implants. Our MAC protocol addresses special communication needs of medical implants like latency, emergency messaging, priority etc., while maintaining an extremely low power-consumption profile. The paper presents the design choices made for a practical cardiac intra-body network and exploits the inherent asymmetries of the network to reduce power consumption. We present a new scheme for deriving analytically the power-optimised TDMA frame parameters like beacon interval and discuss a hardware solution to manage synchronisation overhead. Equations for deriving the duty-cycling efficiency are presented and the packet error rate is calculated for the in-body wireless channel. Our results and simulations show that our protocol is several times more efficient than the state of the art ultra low power protocols. Thus, we illustrate and validate our solution for a very real use case: cardiac networks. However, our new methodology can be applied for any Body Area Network. In this sense, our paper presents a 'universal' solution.

Keywords: Biomedical implants, Body area networks, IEEE 802.15.6, Media Access Layer, Power optimization.

1 Introduction

A body area network (BAN) is a network of sensor nodes which are either on-body or implanted inside the body or both. Since BANs operate in the vicinity of human body, and they have to last for several years in human body, they have a distinct set of requirements of their own which cannot be addressed by the wireless sensor networks (WSN) domain. In view of these requirements, the IEEE has setup a new workgroup (IEEE 802.15.6) which is currently in the process of drafting standards for BANs [1]. The draft intends to propose the medium access layer and physical layer standards for BAN. Since standardization is done for a generic set of requirements, it is not usually optimal for a particular network. In this paper we discuss one such implanted network, an in-body cardiac network and present the design of media access layer for the same. We present the first principles of such a design and justify the design

K.S. Nikita et al. (Eds.): MobiHealth 2011, LNICST 83, pp. 9–15, 2012.

choices thereof. We present a new approach of designing a TDMA-based protocol, for nodes which have different data rates, priorities and latencies. In doing so, we developed an analytical model to determine the trade-off between latency and power consumption, the key issues for medical devices. The paper is organized as follows. In section 2, we present the distinguishing features of an implanted network and discuss our cardiac network and its requirements. In section 3 we present the design of our MAC and its methodology and in section 4 we present the results of our MAC.

2 Requirements of Implanted BANs and Our Use Case

Inherent in the nature of medical implants are requirements and features which make them distinct from other networks. Since these implants deal with critical medical data and have some characteristic communication needs, it is imperative that we understand them before designing the network. Table 1 summarizes the characteristics of the implanted BAN. Having discussed the general requirements of all BANs, we now present the specific use case we use to illustrate our MAC design methodology: a cardiac network. Our cardiac network consists of six different sensors, subcutaneously implanted inside the human body. These sensors are controlled by the pacemaker which functions as the master node placed 15-20 cm away from the nodes. The pacemaker is generally placed just underneath the chest skin. Table 2 specifies the corresponding priorities and data rates of these sensors.

Table 1. Requirements of an implanted BAN

Parameters	Choices
Topology	Star
Nature of traffic	Uplink (mostly)
Power	Ultra low power
Latency	Low and predictable
Priority	Needed
Asymmetry	Between master and slave

Table 2. Specifications of cardiac-BAN

Sensors	Data rates	Priority
PEA	10kbps	High
EGM	5kbps	High
G2D	2kbps	Low
BioZ	1.28 kbps	Low
MV	80bits/second	Low
Temperature	.2 bits/second	Low

3 Design of the New MAC for Cardiac BAN

Given the design choices, we choose TDMA as the access scheme since TDMA eliminate*s collisions, idle listening, overhearing* the major sources of power consumption [3]. Moreover TDMA is the only scheme that can *guarantee* a predictable QoS, a much needed feature for delay bound medical networks *[2].* Despite the aforementioned advantages, TDMA suffers from the fact that the periodic *synchronization* phase must be performed every frame to keep the nodes *synchronized* according to their slots[3]. The other disadvantage of TDMA is that *since nodes* get to speak each frame, the latency of the frame is determined by the frame period. This could be of crucial importance to BANs.

The key parameter in designing a TDMA based MAC is the duty cycling interval (the interval between 2 transmission slots of a node). This interval not only controls latency, but also has a direct effect on the power consumption. In order to conserve power we would like to reduce the sleep period (ON time of radio) as much as possible and then transmit rapidly during our slot and sleep again. However, in BAN unlike the WSN, sensor events are periodic. Therefore the longer the node's sleep, the more data it has in the buffer. It would then have to wake up for a proportionately longer time to send the buffered data to the master. Hence duty cycling the node to reduce power consumption would have only a finite advantage. This implies that if we consider power consumption to be most crucial parameter of optimisation, *there would be a point beyond which sleeping leads to no advantage in power but leads to increase in latency*. We now proceed to determine this point analytically and then present the results for each sensor node of our use case.

Let us consider a duty-cycled system having a current consumption of I_{on} and I_{sl} for on and sleep times respectively. We define t_{sl}, t_{su} and t_{trx} as the radio sleep, start-up and transmission times respectively. Then, the average current drawn over the duty cycling period would be:

$$I_{avg} = [(t_{su}+t_{trx})*I_{on}+ I_{sl}*t_{sl}] / (t_{su}+t_{trx}+t_{sl}) \qquad (1)$$

We define 'R' to be the sampling data rate (in bits-per-second, bps) and 'DR' to be the data-rate over physical layer. So, after sleeping for t_{sl}, the amount of data to be sent and the time to send it are:

$$\text{Data to send} = (t_{su}+t_{sl})* \text{R bits, R is the data sampling rate;} \qquad (2)$$

$$\text{Time to send the data} = t_{trx} = (t_{su}+t_{sl})*(R/DR) \qquad (3)$$

Substituting (3) in (1) we get:

$$I_{avg} = \frac{[t_{su} + (t_{su} + t_{sl}) \times (R/DR)] \times I_{on} + I_{sl} \times t_{sl}}{(t_{su} + (t_{su} + t_{sl}) * (R/DR) + t_{sl})} \qquad (3)$$

We see that for high-rate sensors (high 'R'), the time taken to send the data increases as sleep time t_{sl} increases. Hence the duty cycling efficiency of the sensor node is related to the sampling rate of the inherent sensor. Figure 1 show the graphical description of change in I_{avg} with increase in duty cycle period (up to 5s). The different lines correspond to the sensors of our use-case. The values of t_{su}, R, DR, I_{on}, I_{sl} were taken from the data sheets of the TI-Chicpon CC2430, a popular low power radio [4]. Figure 1 shows clearly the finite effects of duty cycling. Table 3 presents the results for easier comprehension. We see that for high-rate sensors, sleeping beyond 5 or 10 seconds does not lead to any significant reduction in consumption. On the other hand, for low-rate sensors (MV and Temp), the energy minimum is achieved around 100s. Hence, we see that each sensor has a different energy minimum point.

In table 4, we have underlined the ideal duty cycle points for each sensor. For example, for the PEA sensor, going from a 5s beacon interval to 10s leads to reduction of only 3μA. If we can tolerate 3μA of additional consumption, we can reduce the latency from 10s to 5s. We could also duty-cycle the PEA at 1s intervals but the energy cost would be 30μA higher. This table therefore takes us back to the latency-versus-energy trade-off that we have just discussed. *Furthermore, the duty-cycle points naturally lead us to the ideal beacon intervals for each of these sensors in the TDMA scheme*. Now, how do we incorporate these different ideal beacon intervals of each sensor in the TDMA scheme? We propose to do so by choosing the beacon interval such that it meets the latency and power requirements for the higher- rate and priority sensors.

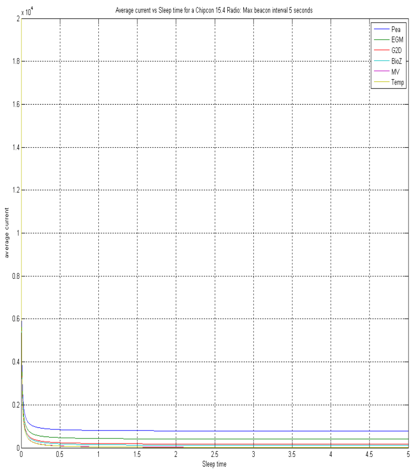

Fig. 1. Decrease in average current as duty cycle increases for all sensors

Table 3. Average current consumption(μA) of our sensors for different duty cycle intervals

Sensors	.1s	.5s	1s	5s	10s	100s
PEA	1147	847	808	777.9	774	770
EGM	777	471	432	401	397	393
G2D	548	238	198	170	163	160
BioZ	493	182	142	110.8	106	103
MV	399	87	47	15.39	11	7.99
Temp	393	80	40	8.997	4.99	1.4

Table 4. Cost of Synchronizing for different sensors

Snsors	Oringinal Iavg (μA)	Iavg with early start(μA)	Additional Cost(μA)
PEA:	777.9	793.5	15.6μA
EGM:	401	417.2	16.2μA
G2D:	167	184.3	17.3μA
BioZ:	110.8	127.6	16.8μA
MV:	7.99	9.4	1.41μA
Temp:	1.4	3.0	1.6μA

In our case, the high-rate sensors are also those that have higher priority: EGM and PEA. We thus choose a beacon interval according to these two sensors, and then make the other sensors wake up after every 'N' beacon intervals. We choose a beacon interval of 5s, the MV and Temp sensors, which have the ideal beacon interval of 100s, wake up after every 20 (100/5) beacons. Thus, not all nodes wake up with every beacon: Nodes which have more data to transmit (PEA/EGM) wake up every 5s; nodes with less data to send (G2D/BioZ) wake up after 10s; and so on. Therefore the TDMA scheme devised by us does not mandate that each sensor wake up with each beacon. Having decided the *beacon interval* we now discuss the crucial problem of *synchronisation* of sensors.

The more the nodes sleep, the more time drift they accumulate. So, nodes that wake up after 20 beacons (100s) would have much higher time uncertainty about their time slots and could instead transmit in their neighbouring slot. We propose to resolve this problem as follows: If we have a crystal of tolerance 'ε' ppm, the amount of timing accuracy 'δ' over the duration of sleep interval (same as the beacon interval BI) will be($\pm \varepsilon \times BI$) . Since we analytically know ε, and the maximum timing error that we can encounter (δ), we can avoid missing the beacon by forcing the radio start-up at (BI-δ) instead of BI. We call this the '*cost of synchronising*' with the beacon. This cost will be higher for nodes which sleep longer. Table 4 shows the cost of synchronising with beacon in terms of current for different sensors for the minimum-energy points of table 4. Note that for low-rate sensors, the relative cost is insignificant (2μA). For the higher-rate sensors, it is low.

Slots Determination

Table 5 shows the time required by each sensor to do so. We have assumed conservative values of PER, overhead and retransmission to arrive at the worst-case scenario. We chose the slot interval to be 20ms, since 20ms was close to the least.

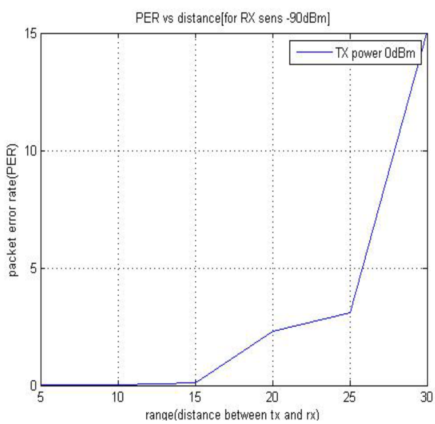

Fig. 2. PER versus distance simulation

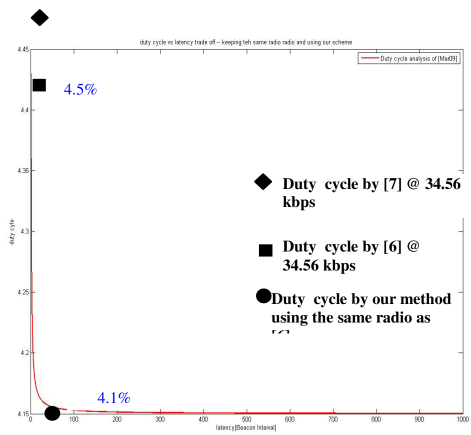

Fig. 3. Gain in duty cycle from 4.51% [6] to 4.1%

Table 5. Slot size determiniation by considering various factors

Sensors	BI	Total Data @BI (BI*DR)	Time To Send (@250kbps)	+ Over -head (15%)	+ PER (@5%)	+ACK(20%)	#of slots (20 ms slots)
PEA	5s	50kb	200ms	230ms	242ms	290ms	15
EGM	5s	25kb	100ms	115ms	121ms	145ms	8
G2D	10s	20kb	80ms	92ms	97ms	116ms	6
BioZ	10s	12.8kb	51.2ms	59ms	62.5ms	75ms	4
MV	100s	8kb	32ms	37ms	39ms	47ms	3
Temp	100s	.02kb	.08ms	.108ms	.11ms	.13ms	1

4 Results

We simulated the network in the network simulation software OMNET++, a popular discrete event network simulator [5]. The MICS channel model as specified by the IEEE 802.15.6 was used to model physical layer behaviour [1]. We carried out a packet error rate (PER) analysis of our protocol inside a slot. For a receiver sensitivity of -90dBm, we tried to find out PER behaviour for various distances [4]. This gives us a good idea of the range of our network. Figure 6 shows the plot of PER as we vary the range for MICS band. We obtain an acceptable PER (3.1%) for distances upto 25cm inside body. We now compare our work with the two ultra-low power TDMA-based BAN protocols by Omeni [7] and Marinkovic [6]. Note that a power analysis depends on the underlying hardware, hence radios which are more power-efficient and have faster data rates tends to give better power consumptions. The radios used by [6] and [7] have much lower data rates (34.56kbps and 50kbps respectively) compared to our solution (250kbps). A direct analysis on the basis of duty cycle or power consumption would automatically favour our implementation. So, we try to provide two analyses by: a) Comparing our scheme “as it is” with these two protocols, b) Adapting our cycle scheme to use the physical radio characteristics of these protocols. Duty-cycle analysis is generally considered a good figure of merit for any

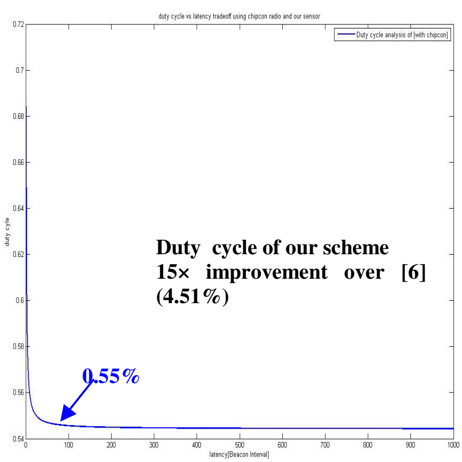

Fig. 4. Duty cycle analysis for a 2.5kbps sensor radio

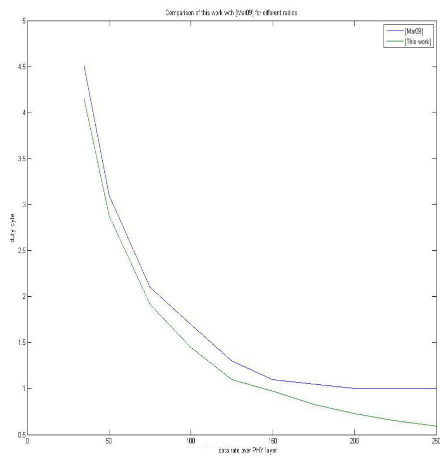

Fig. 5. Comparison of our scheme using our scheme with [6] for different radios

TDMA protocol [6]. It measures how much time one’s receiver is ‘on’. The protocol published in [6] reports a duty cycle of of 4.51% for 1.25kbps sensor and 5.7% for 2.5kbps. We carried out this analysis for the protocol of [6] and [7] and found that using the same physical and MAC parameters and simply altering the beacon interval, we can reduce the duty cycle from 4.51% to 4.10%, as shown in figure 3. This is the reduction for each frame; the overall gains in the energy over the life-time of sensor would be significant. Using our radio and our MAC scheme, the same data-rate sensor can be duty-cycled to 0.55%, a factor of more than 9 times improvement over and 15.5 times over [6](figure 4) . Figure 5 shows the duty-cycle analysis as we change the data rate over physical layer (use faster radios). We see that our scheme shows greater gains as we move to faster and better radios. These results assume more importance because of the fact that most commercial low-power radios have higher data rates (>200kbps)[4],[1]. Furthermore, our protocol takes into account retransmission while other protocols permit some packet loss which could be critical for medical data.

References

1. https://mentor.ieee.org/802.15/documents?is_group=0006
2. Keshav: Engineering Approach to Computer Networking. An: ATM Networks, the Internet, and the Telephone Network (1996)
3. Ye, W., Heidemann, J., Estrin, D.: An energy-efficient MAC protocol for wireless sensor networks. In: Proc. IEEE 21st Ann. Joint Conf. IEEE Comput. Commun. Soc. (2002)
4. Texas instruments, Chipcon CC2430 data sheet
5. http://www.omentpp.org
6. Marinkovic, S.J., et al.: Energy-Efficient Low Duty Cycle MAC Protocol for Wireless Body Area Networks. IEEE Transs. on Information Technology in Biomedicine (2009)
7. Omeni, O., et al.: Energy Efficient Medium Access Protocol for wireless medical body area sensor networks. IEEE Trans. on Biomedical Circuits and Systems (December 2008)