

An Energy Analysis of IEEE 802.15.6 Scheduled Access Modes for Medical Applications

Christos Tachtatzis, Fabio Di Franco, David C. Tracey,
Nick F. Timmons, Jim Morrison

WiSAR Lab,
Letterkenny Institute of Technology,
Port Road, Letterkenny, Co. Donegal, Ireland
{christos,fabio,david,nick,jim}@wisar.org

Abstract. Medical body area networks will employ a range of implantable and body worn devices to support a wide range of applications with diverse QoS requirements. The IEEE 802.15.6 working group is developing a communications standard for low power devices operating on, in and around the body and medical devices are a key application area of the standard. The ISO/IEEE 11073 standard addresses medical device interoperability and specifies the required QoS for medical applications. This paper investigates the lifetime of devices using the scheduled access modes proposed by IEEE 802.15.6, while satisfying the throughput and latency constraints of the ISO/IEEE 11073 applications. It computes the optimum superframe structure and number of superframes that the device can sleep to achieve maximum lifetime. The results quantify the maximum expected achievable lifetime for these applications and show that scheduled access mode is not appropriate for all application classes such as those with intermittent transfer patterns.

Key words: Energy Analysis, IEEE 11073, IEEE 802.15.6, Scheduled Allocations, Wireless Body Area Network, Wireless Medical Applications

1 Introduction

The rapid expansion of wireless technology has led to the possibility of widespread untethered medical and health monitoring. The use of wireless technology, promises benefits in terms of replacing cabling, greater flexibility in equipment placement, wider access to patient data (not limited to the bedside or wired points), patient mobility in hospital and possibly home monitoring allowing earlier patient release. There will also be opportunities for the emerging monitoring and alerting applications such as remote patient monitoring and automatic drug delivery.

The emphasis in Wireless Sensor Networks (WSNs) and Wireless Personal Area Networks (WPANs) has been low power, low cost and short range operations. The importance of low power operation is even greater in medical Wireless Body Area Networks (medical WBANs) where devices are expected to operate over long periods without battery replacement and charging may not be feasible;

i.e. implantable medical devices. Existing WPAN standards have high energy demands for medical application and insufficient QoS guarantees. IEEE 802.15.6 [1] has the potential to overcome the limitations of other standards, such as IEEE 802.15.1 [2] and 802.15.4 [3] and to allow the wider implementation and deployment of Wireless Body Area Networks (WBANs). IEEE 802.15.1 was the first standard to focus on the short range personal area networking environment and IEEE 802.15.4 traded throughput for low power operation. Previous work [4] and [5] has shown lifetime limitation of IEEE 802.15.4 for medical applications.

Recent work has considered the performance of the IEEE 802.15.6 contention based access under saturation conditions [6]. The delay and throughput limits for a single device using contention based access was examined in [7]. Both [6] and [7] considered contention based access and in this paper we extend this by examining the performance of the contention free scheduled access modes. The results presented are based on the analysis in [8]. In particular the medical device lifetimes are investigated for real application requirements from the ISO/IEEE 11073 [9]. These results are computed with a mixed integer program that finds the optimum superframe structure and the number of beacon periods through which the device could sleep to achieve maximum lifetime, while satisfying the QoS requirements (data rate and latency) for a range of applications.

The remainder of the paper is organised as follows: Section 2 gives an overview of IEEE 802.15.6. Section 3 gives an overview of the ISO/IEEE 11073 application scenarios. Section 4 presents the analysis and the mixed integer program. Sections 5.1 and 5.2 presents the sensitivity of the lifetime to the MAC parameters and the lifetime estimates respectively. Section 6 presents future work and conclusions.

2 IEEE 802.15.6 Overview

2.1 Physical Layer

IEEE 802.15.6 specification defines Narrowband (NB), Ultra-Wide Band (UWB) and Human Body Communications (HBC) physical layers and a common frame structure. The NB physical layer operates in seven different frequency bands with a variable number of channels, bit rates and modulation schemes. Our analysis concentrates on bands 6 and 7, which operate at symbol rates of 600 ksp/s with varying modulation schemes, coding rates and spreading factors.

2.2 Medium Access Control Layer

The IEEE 802.15.6 draft specifies a common MAC for all the supported physical layers and which can use one-hop star or two-hop restricted tree topologies. In these topologies, the hub is responsible for coordinating channel access by establishing one of the following three access modes:

- Beacon mode with beacon period superframe boundaries

- Non-beacon mode with superframe boundaries
- Non-beacon mode without superframe boundaries

The time base is divided into equal length beacon periods (also known as superframes) and each superframe is divided into allocation slots. In the first two access modes the time base is common between hubs and nodes and is decided by the hub; i.e. the hub establishes superframe boundaries and defines the number of allocation slots in it. In the first access mode, the hub communicates the superframe structure via beacon frames or Timed frames (T-Poll). The second access mode does not transmit beacons and the superframe structure is enforced through the use of Timed frames (T-Poll). In the non-beacon mode without superframe boundaries each node establishes its own time base independently.

In the beacon mode with superframe structure the hub organises the superframe in access phases, shown in Fig. 1, and allows three types of access:

1. Random access (Contention based): CSMA/CA or Slotted ALOHA for the narrowband and ultra-wide band physical layers respectively. These are the EAP1, EAP2, RAP1, RAP2 and CAP shown in Fig. 1. The EAPs are reserved for emergency high priority traffic while the RAPs are used for non-recurring transfers.
2. Improvised, unscheduled access: Post (i.e. a hub instruction) or Poll (i.e. a data request from the hub). During this mode, devices must be awake and wait for a poll or post frame from the hub, before they can transmit.
3. Scheduled access (Contention free): 1-periodic where devices exchange frames with the hub in every superframe or m -periodic where devices and hubs exchange frames every m superframes allowing the device to sleep between transfers. In this mode, devices can start their transfer when the reserved allocation slot time has commenced

Scheduled transfers, unscheduled and improvised transfers occur in the Managed Access Phases MAP1 and MAP2.




Fig. 1. Beacon mode with beacon period superframe boundaries structure.

The length of these phases is variable and is defined in numbers of allocation slots. The draft defines four acknowledgement policies: 1) not acknowledged frames (N-Ack), 2) immediately acknowledged (I-Ack), 3) block acknowledged later (L-Ack) and 4) block acknowledged (B-Ack).

3 ISO/IEEE 11073 Applications

The ISO/IEEE 11073 Draft for Point-of-Care (PoC) medical devices [9] defines a range of medical application classes. Table 1 is a subset of the IEEE 11073 application classes which can be supported by the maximum data rate of 971 kbps provided by the narrowband physical layer of the IEEE 802.15.6 draft. The ultra-wideband physical layer of IEEE 802.15.6 can support application data rates up to 10 Mbps.

Class: Data Type	Latency	Bandwidth
A: Alarms/alerts/ Positional Alerts (real-time)	A1: < 200ms and A2: <3 s	64 bytes per alarm
B: Patient State	< 3 s	64 bytes per alarm
C: Sensor watchdog/heartbeat	< 60 s	64 bytes per hour
D: Reminder	< 3 s	1632 bytes per alarm
E: Physiologic parameters (real-time) [e.g. episodic BP, HR, SpO2, ETCO2, temp]	< 3 s	20 bytes/param at E1: 0.5 to E2: 5 Hz
F: Telemetry Waveforms (real-time)	< 300ms	ECG: [F1: 3-lead 2.4 kbps, F2: 5-lead 10 kbps, F3: 12-lead 72 kbps], F4: Ventilator: 50-60 bps, F5: SpO2: 50-120 bps

Table 1. Applications from IEEE 11073.

4 Analysis

The analysis presented is for medical monitoring devices which report periodically. As such it concentrates on IEEE 802.15.6 Scheduled Access mode (1-periodic and m-periodic allocations in beacon mode with superframe boundaries) using the one-hop star topology. The device lifetimes are determined using the analytical model presented in [8] for uplink block transfers¹. The purpose of this analysis is to determine the superframe structure and the number of beacon periods through which the device must sleep to achieve maximum lifetime for these applications. It considers the end device in a one-hop point-to-point link. This

¹ The analytical model was developed for version 1 of the draft and it was updated to comply with version 2. See discussion further in this section.

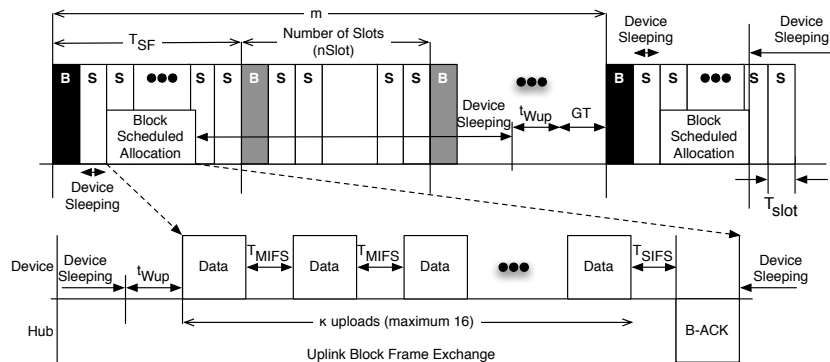


Fig. 2. Optimisation parameters. Superframe structure and periodicity m .

is the best case for a particular device as there are no conflicting requirements from other devices on the BAN.

The IEEE 802.15.6 offers a flexible superframe structure that can be adjusted by the hub to suit the communication requirements of the BAN. Fig. 2 shows the superframe structure and the device’s duty cycle for an m -periodic uplink allocation. The cycle starts with the device receiving a beacon and subsequently going to sleep until the start of its allocation interval where the device wakes-up for the block scheduled allocation. After the block schedule allocation the device goes to sleep for m superframes before waking-up to receive the next beacon and the cycle repeats itself. During the block schedule allocation the device sends data to the hub in block acknowledged transfer. The maximum number of frames in a block transfer is 16 according to the IEEE 802.15.6. Each frame fragment has the maximum size of 255 bytes except for the last frame fragment which has the length required to send the remaining data. This analysis uses the maximum frame size, because it is the most efficient transfer due to the reduced header overhead under ideal channel conditions, which allows us to determine the upper bound of maximum lifetime. It may not be suitable in scenarios where the bit error rate is high and there will be a higher probability of frame corruption.

If the data to be transferred cannot fit in a single scheduled allocation, the device requests additional allocations in the same superframe and it is assumed that the device goes back to sleep between the allocations intervals.

The superframe is divided in slots and the number of slots ($nSlots$) is adjustable by the hub. To allow further flexibility the hub can adjust the duration of the slot in increments of $pAllocationSlotResolution$ (equal to 1 ms for the narrowband physical layer). The minimum slot duration $pAllocationSlotMin$ is defined equal to 1 ms for the narrowband physical layer.

The superframe duration is:

$$\begin{aligned}
 T_{SF} &= nSlots \cdot T_{slot} \\
 T_{slot} &= (pAllocationSlotMin + L \cdot pAllocationSlotResolution) \quad (1)
 \end{aligned}$$

The parameters that control the superframe structure are $nSlots$, which can be between 1 and 256 and the slot length L , which can be between 0 and 255. These parameters ranges allow slot durations between 1 ms and 256 ms and superframe durations in the range of 1 ms to 65.536 s.

The analysis determines the optimal value of m for maximum device lifetime when the m-periodic mode is used and observing the constraints specified by IEEE 802.15.6 and radio constraints such as device warm-up times [8]. The guard time used in the analytical model of [8] has been updated to comply with the definition of [1] as follows.

The nominal synchronisation interval (SI_n) is specified to be 8 beacon period lengths (T_{SF}). The nominal guard time is $GT_n = GT_0 + 2 \cdot SI_n \cdot \delta$ where GT_0 is fixed at 61.6 μ s based on data-link and physical layer parameters, while δ is the clock accuracy. GT_n should be used when the last synchronisation interval SI is less than the nominal SI_n , otherwise GT_a should be appended to it. Hence the guard time is a function of SI and after rounding to the clock accuracy (δ):

$$GT(SI) = \begin{cases} \lceil \frac{GT_n}{\delta} \rceil \cdot \delta & \text{if } SI < SI_n \\ \lceil \frac{GT_n + GT_a}{\delta} \rceil \cdot \delta & \text{if } SI \geq SI_n \end{cases} \quad (2)$$

The optimisation problem was formulated as a mixed integer program that maximises the device lifetime, assuming that no other devices operate on the BAN and finds the optimum parameters that maximise the following cost function:

$$\begin{aligned} \text{maximise :} & \quad T_{life} = \frac{Q}{I_{total}(nSlots, L, m, \Delta, \tau)} \\ \text{where :} & \quad \Delta = \frac{N_F \cdot 255 \cdot 8}{m \cdot T_{SF}}, \tau = m \cdot T_{SF} \\ \text{subject to :} & \quad \Delta \leq \Delta_{App}, \tau \leq \tau_{App} \\ & \quad nSlots \in [1, 256], L \in [0, 255], m \in [1, 256] \end{aligned} \quad (3)$$

where T_{life} is the device lifetime in years, Q the battery capacity, Δ the achieved data rate, τ the achieved delay, I_{total} the average current consumption of the device as a function of ($nSlots, L, m, \Delta, \tau$) and is determined as in [8], N_F the number of data frames transmitted every m superframes and T_{SF} the superframe duration. The application data rate and latency requirements are Δ_{App} and τ_{App} respectively.

5 Device Lifetime For Medical Applications

Equation 3 is solved for selected data rates and latency requirements per IEEE 11073. The solutions give the superframe structure ($nSlots, L$) and the parameter m for m-periodic allocations that maximise the device lifetime. It is important to note that more than one equivalent solution may exist. For example, for the

nSlots	L	T_{SF}	m	Wake-up Period (s)	Lifetime (years)
15	0	15	20	0.3	0.14411
4	4	20	15		
5	3				
10	1				
20	0				
5	4	25	12		
25	0				
2	14	30	10		
3	9				
5	5				
6	4				
10	2				
15	1				
30	0				

Table 2. Solutions for ECG 5-Lead, Data Rate 10 Kbps, Latency 0.3 s.

class F2: 5-lead ECG from Table 1 with data rate of 10 kbps and latency requirement of 300 ms there are fourteen optimum equivalent solutions as illustrated in Table 2.

For the application class shown and the other classes considered, the equivalent solutions for maximum lifetime are dominated by their latency requirement. All the parameter combinations that maximise the wake-up period result in equivalent lifetimes. However, there are a few notable exceptions. For example, solutions where the parameter m is less than 8 are sub-optimal. This is explained by the fact that the nominal guard time as specified in the draft standard is fixed and over-provisioned when the last synchronisation interval (SI) is below the nominal synchronisation interval (SI_n). In fact, the version 1 of the draft had specified the nominal guard time to be equal to 1/10 of the allocation slot (T_{slot}) and made any solution with $L > 0$ suboptimal in terms of energy efficiency, making the parameter L meaningless. The version 2 of the draft has revised the definition of nominal guard time to be proportional to the beacon transmission frequency, making L a parameter worth controlling in optimising the lifetime. This modification to the draft has, however, resulted in a similar limitation on optimising the lifetime using parameter m , i.e. optimal lifetimes have m greater than 8, unless the application's data rate forces m to be less.

Note that there are no solutions with small superframe duration (T_{SF}) in Table 2. Such solutions are not feasible because the duration of the superframe is too short for the block transfer to fit between the transmissions of two beacons. For example, a superframe with two slots and L equal to zero will result in a superframe with duration of 2 ms and the parameter m to maximise the wake-up period must be 150. In this case the duration of the block transfer is 12.69 ms and can not fit within the 2 ms superframe. The mixed integer program uses constraints to remove such solutions (not shown in Equation 3 for clarity).

5.1 Device Lifetime Sensitivity to MAC parameters

The dynamics and sensitivity of the model can be seen in Figures 3 to 5 for selected application classes (classes A and B not shown in Figures, and the results are similar to class E). The maximum lifetime against the parameter L , which is proportional to the slot length, is shown in Fig. 3. For each point in this graph the parameter L is kept constant at the values indicated on the x-axis and the other parameters (number of slots and m), are allowed to change in the mixed-integer program. The output of the integer program returns the maximum achieved lifetime for the specified value of L . The circles in the graph indicate the maximum points for each application class (i.e. the maximum of the maximum lifetimes for the given L).

The step change in maximum lifetime at certain values of L (Fig. 3), must be noted and is an outcome of meeting the latency requirements of an application for a given L . For example, the 12-lead ECG requires a latency of 300 ms. A value of $L=99$, (so $T_{slot}=100$ ms) meets this latency with 3 superframes and 1 slot in each superframe. When L is increased to 100 (so T_{slot} is 101 ms), and holding the number of slots in the superframe at 1, leads m having to decrease to 2 in order to satisfy the latency requirement. This results in a more frequent duty cycle of 202 ms which increases current consumption and reduces the lifetime.

The maximum lifetime against the number of slots in a superframe is shown in Fig. 4. From the graph it can be observed that multiple equivalent solutions




Fig. 3. Maximum lifetime versus slot length (L) for applications C, D, E1, E2, F1 and F2.

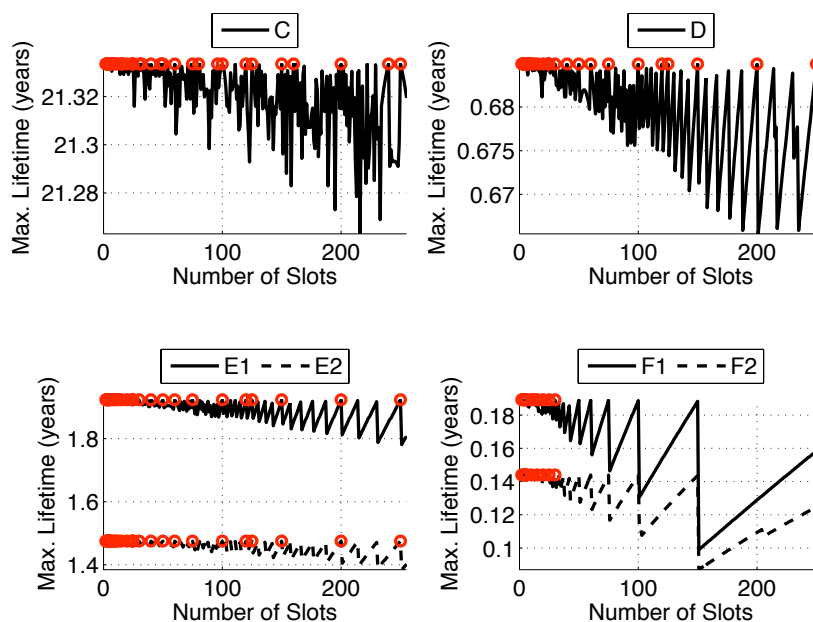


Fig. 4. Maximum lifetime versus number of slots in a superframe ($nSlots$) for applications C, D, E1, E2, F1 and F2.

exist for each application. In a similar fashion to Fig. 3, a step change can be observed at certain values of $nSlot$ caused by meeting the enforced application latency requirement. The circled points in Fig. 4 show the equivalent solutions that maximize the maximum device lifetime.

The maximum device lifetime can be plotted against the parameter m as shown in Fig. 5. It must be pointed out that not all applications have feasible solutions for $m \in [1, 255]$, because high values of m result in more data building up in the device buffers and so require longer superframes for the data transfer. This combination of longer superframes and high m will violate the latency constraint of the applications, e.g. application F1: 3-lead ECG is only feasible when m is in the range of 1 to 21.




Fig. 5. Maximum lifetime versus periodicity m for applications C, D, E1, E2, F1 and F2.

5.2 Device Lifetime Estimates

Fig. 6 shows the maximum achieved lifetime for variable data rates and latency requirements. From Fig. 6 it can be observed that lifetime greater than 1 year can only be achieved by applications with a latency constraint of 3 s or more.

The maximum device lifetime for applications with latency constraints of 0.2 and 0.3 s, and data rates less than 1000 bps is 51 and 76 days respectively. This low device lifetime can be explained by the fact that the superframe duration is kept low to satisfy the tight delay constraint causing devices to wake up regularly to send small chunks of data and receive beacons.

The frequent wake up cycle, results in the consumption of a significant amount of energy in beacons. This is validated in Fig. 7, which shows the percentage of energy spent in beacons. For low latency constraints and low data rates more than half of the device's energy is spent in handling beacons.

Fig. 6 also shows that at high data rates, the proportion of energy spent in transferring data dominates as illustrated by the convergence of the lifetimes for different values of application latency. This dominance is also evident in Fig. 7 where the percentage of energy consumption associated with beacons is below 10% at high data rates.

From Fig. 6 the maximum lifetime for the applications classes of Table 1 can be extracted and are summarised in Table 3. Table 3 also shows one of the computed optimum solutions (the one with the maximum number of slots




Fig. 6. Maximum device lifetime satisfying the given data rate and latency requirements.




Fig. 7. Normalized energy spent in beacons.

in the superframe) for each application from ISO/IEEE 11073. The selected solutions shown are those with the maximum number of slots in the superframe. The additional slots will enable the hub to accommodate connections with other devices with the extra slots that remain free.

While it was expected that the scheduled access modes would not be suitable for intermittent data transfer, because devices with m -periodic allocations have to wake-up to receive beacons and maintain their slot allocation even if they have no data to transfer, Table 3 shows how low the lifetime would be in such cases. Use of scheduled access modes for classes such as A and D, could only be justified when high reliability and deterministic response time is required at the expense of energy usage. For class A, use of the Emergency Access Phases (EAPs) would be more appropriate, while class D could use the Random Access Phases (RAPs).

Finally, the solution for application class F3: 12-lead ECG, with data rate of 72 kbps and delay constraint of 300 ms, is only feasible when the parameter m is less or equal to 6. This is caused by the high data rate demand of the application. For m more than 6, too much data are accumulated in the device buffers and the superframe length does not fit the data exchange forcing m to reduce.

App. Class	App. Rate	App. Latency	nSlots	L	m	Lifetime (years)
A1	2048	0.2	25	0	8	0.13095
A2	2048	0.3	30	0	10	0.19156
B	170.6	0.3	30	0	10	0.20489
C	0.12	60	250	29	8	21.3336
D	4352	3	250	0	12	0.68493
E1	80	3	250	0	12	1.92350
E2	800	3	250	0	12	1.47530
F1	2400	0.3	30	0	10	0.18928
F2	10000	0.3	30	0	10	0.14411
F3	72000	0.3	50	0	6	0.04839
F4	60	0.3	30	0	10	0.20581
F5	120	0.3	30	0	10	0.20522

Table 3. One optimum solutions for ISO/IEEE 11073 applications of Table 1.

6 Conclusions

Medical Devices in BANs are an important application area for IEEE802.15.6 and maximising device lifetime is a key requirement in such scenarios. The analysis and results presented provide best case estimates for device lifetime when

using IEEE 802.15.6 scheduled access modes with data rate and latency constraints as defined by ISO/IEEE 11073 for medical devices. The paper has shown the optimum superframe structure and m-period for these application scenarios, which were found using a mixed integer program.

The functionality in IEEE 802.15.6; to allow devices to skip beacon periods provides flexibility and reduces energy consumption. This is, however, subject to the latency requirements of given applications which may force devices to wake-up more frequently and limits the amount of time the devices can spend in sleep state. The results presented have shown the significant extent to which these application constraints, particularly latency, can impact the device lifetime, when scheduled access modes are used.

The findings on the estimated device lifetimes (and corresponding superframe structure), such as the low device lifetime for applications with intermittent data transfer, show the importance of considering medical application requirements to select the appropriate access mode and ensure the best use of the proposed standard.

Analysis of the first draft, which specified the nominal guard time to be equal to 1/10 of the allocation slot, made any solution with $L > 0$ suboptimal in terms of energy efficiency. Draft 2 redefined the nominal guard time to be fixed (and proportional to the nominal synchronisation interval), resulting in the optimum lifetime solutions having the device sleep for more than 8 beacon periods, unless the application's data rate forces m to be less. This use of a fixed nominal guard time to mitigate the effects of missing beacons should be reconsidered, given the impact it has on the device lifetime.

Future work will be necessary to characterise random, improvised and unscheduled access modes and subsequently investigate the scenario where multiple applications with contradictory QoS constraints operate in the same body area network. This would provide data to form policies based on application requirements and aid the hub to set the MAC parameters and achieve desired outcomes, i.e. maximum device lifetime.

Acknowledgment

This work was carried out under the auspices of Enterprise Ireland Applied Research Enhancement (ARE).

References

1. IEEE P802.15.6/D02 Wireless Medium Access Control (MAC) and Physical Layer (PHY) Specifications for Wireless Personal Area Networks (WPANs) used in or around a body, December (2010).
2. IEEE Standards for Information Technology - Telecommunications and Information Exchange Between Systems - Local and Metropolitan Area Networks - Specific Requirements - Part 15.1: Wireless Medium Access Control (MAC) and Physical

- Layer (PHY) Specifications for Wireless Personal Area Networks (WPANs), IEEE Std 802.15.1-2005 (Revision of IEEE Std 802.15.1-2002), pp. 1-580, (2005)
3. IEEE Standards for Information Technology - Telecommunications and Information Exchange Between Systems - Local and Metropolitan Area Networks - Specific Requirements - Part 15.4: Wireless Medium Access Control (MAC) and Physical Layer (PHY) Specifications for Low Rate Wireless Personal Area Networks (WPANs), IEEE Std 802.15.4-2006 (Revision of IEEE Std 802.15.1-2003), pp. 1-305, (2006)
 4. Timmons, N.F., Scanlon, W.G.: Analysis of the performance of IEEE 802.15.4 for medical sensor body area networking. 1st Annual IEEE Communications Society Conference on Sensor and Ad Hoc Communication Networks (IEEE SECON), October (2004)
 5. Sukor, M., Ariffin, S., Faisal, N., Yusof, S., Abdallah, A.: Performance study of wireless body area network in medical environment. 2nd Asia International Conference on Modeling and Simulation (AICMS), May (2008)
 6. Rashwand S., Mistic, J., Khazaei, H.: Performance Analysis of IEEE 802.15.6 Under Saturation Condition and Error-Prone Channel. IEEE Wireless Communications and Networking Conference, March, (2011)
 7. Ullah S., Kwak, K. S., Throughput and Delay Limits of IEEE 802.15.6, IEEE Wireless Communications and Networking Conference, March, (2011)
 8. Tachtatzis C., Di Franco, F., Tracey, D., Timmons, N.F., Morrison, J.: An energy analysis of the IEEE 802.15.6 scheduled access modes. IEEE Globecom Workshop on Mobile Computing and Emerging Communication Networks (MCECN), December, (2010)
 9. Draft Health Informatics - Point-of-Care Medical Device Communication - Technical Report - Guidelines for the use of RF wireless technology, IEEE Unapproved draft Std P11073-00101/D5, June, (2008).