Overview of the IEEE 802.15.4/4a standards for low data rate Wireless Personal Data Networks

Luca De Nardis and Maria-Gabriella Di Benedetto
Infocom Department
School of Engineering
University of Rome La Sapienza
Rome, Italy 00184
Emails: {lucadn, dibenedetto}@newyork.ing.uniroma1.it

***Abstract*— The IEEE 802.15.4 standard provides a framework for low data rate communications systems, typically sensor networks. The 802.15.4a revision introduces new options for the physical layer, in order to support higher data rates and accurate ranging capability, enabling new applications based on information on distance and positions of the devices in the network. In this paper the differences among physical layers of 802.15.4 vs. 802.15.4a are briefly reviewed. Next, the MAC layer of 802.15.4 vs. 802.15.4a are reviewed and compared. Device functionalities, network topologies as well as access strategies in both standards are described, and the impact of the new physical layer features on MAC and higher layers are discussed, with particular attention to the ranging scheme adopted in the new revision of the standard.**

I. INTRODUCTION

The goal of the 802.15.4a standard for Low Data Rate (LDR) networks, currently under definition within the IEEE, is to provide joint communications and high accuracy positioning in future sensor networks [1]. It represents therefore an evolved version of the 2003-released parent IEEE 802.15.4 standard, in which the positioning aspect was not addressed [2]. The IEEE 802.15.4a adopts an Impulse Radio Ultra Wide Band (IR-UWB) physical layer, capable of providing accurate ranging and therefore favor accurate positioning, and a low rate chirp-based physical layer.

In a parallel activity, the former IEEE 802.15.4b Task Group focused on refining 802.15.4, by removing ambiguities and addressing issues raised during early implementation efforts of 802.15.4 devices. The 4b Task Group completed its activities in 2006 by releasing an updated version of the original standard [3].

In this paper, we provide an overview and comparison of 802.15.4 vs. 802.15.4a with respect to the MAC layer, taking into account the impact of the characteristics of the new physical layer on medium access. The paper is organized as follows. Section II provides a brief description of the original physical layer defined in 802.15.4 and of the new physical layers defined in 802.15.4a, highlighting the differences among the two standards. Section III describes device functionalities and network topologies. Section IV focuses on access strategies and resource management within a network, describing the CSMA-CA strategy defined in 802.15.4 and the new ALOHA option introduced in 802.15.4a. Section V describes the ranging procedure defined in the 802.15.4a, highlighting its challenges and the potential impact on MAC procedures. Section VI draws conclusions.

II. OVERVIEW OF PHYSICAL LAYERS

The original 802.15.4 standard, released in 2003 [2], adopted a wideband physical layer using a Direct Sequence Spread Spectrum technique (DSSS). The standard provided specifications for operating in three different frequency bands: the 868 MHz band, available in Europe, the 915 MHz band, available in US, and the 2400 MHz ISM band, available worldwide. In order to allow the coexistence of several networks in the same location a Frequency Division Multiplexing (FDM) approach was adopted, by dividing each of the bands in channels; the access strategy within a single network was based on time division, as shown in Section IV. The channelization scheme was the following:

- 1 channel (Channel 0) was defined in the 868 MHz band;
- 10 channels (Channels 1-10) were defined in the 915 MHz band, with a channel spacing of 2 MHz;
- 16 channels (Channels 11-26) were defined in the 2.4 GHz band, with a channel spacing of 5 MHz.

The channels defined in the first two bands were intended for very low bit rate operations, with user rates of 20 kb/s and 40 kb/s per channel for the 868 MHz and the 915 MHz bands respectively. The 2.4 GHz ISM band, on the other hand, allowed bit rates up to 250 kb/s per channel, thanks to the larger bandwdith allocate to each channel. It is worth mentioning that the improved version of the standard released in 2006 introduced new modulation schemes for the channels in the 868 MHz and 915 MHz bands, allowing to achieve 250 kb/s per channel for these two bands as well [3].

Shortly after the release of the 802.15.4 standard it was evident that the range of potential applications for a low bit rate standard could be signficantly increased by the capability of measuring distance between devices in the network with high accuracy. Since this capability was precluded to 802.15.4 devices due to the limited signal bandwidth, the 802.15.4a Task Group was created with the goal of defining a new physical layer, able to provide the desired ranging capability, and correspondingly adapting the Medium Access Control layer.

TABLE I
CHANNELIZATION SCHEME IN IEEE 802.15.4A

Channel	Center frequency (MHz)	Bandwidth (MHz)
0	499.2	499.2
1	3494.4	499.2
2	3993.6	499.2
3	4492.8	499.2
4	3993.6	1331.2
5	6489.6	499.2
6	6988.8	499.2
7	6489.6	1081.6
8	7488.0	499.2
9	7987.2	499.2
10	8486.4	499.2
11	7987.2	499.2
12	8985.6	499.2
13	9484.8	499.2
14	9984.0	499.2
15	9484.8	1355

The 802.15.4a physical layer is based on two different technologies: Ultra Wide Band (UWB) and chirp signals.

The UWB signal adopted in 802.15.4a uses three frequency bands within the range of frequencies made available by regulation for UWB emissions released in 2002 by the Federal Communications Commission (FCC) [4]: a sub-GHz band, a low band in the 3 - 5 GHz range, and a high band in the 6 - 10 GHz range. As in the original physical layer, the bands are divided in channels, as presented in Table I.

A device is not required to support all channels; for each band there is however a mandatory channel that all compliant devices must support (these are channels 0, 3 and 9 for sub-GHz, low and high band, respectively). The standard channels, characterized by a bandwidth of about 500 MHz, are expected to allow ranging with an accuracy in the order of 1 meter. Applications requiring higher ranging accuracy may use channels 4, 7 and 15, which offer a larger bandwidth, and allow thus for a higher ranging accuracy. It should be noted furthermore that up to two 802.154.a networks can operate in the same channel at the same time, thanks to the adoption of preambles, characterized by low cross-correlation. The preambles are defined within the standard.

The UWB physical layer uses an Impulse Radio approach, in which short pulses with a bandwidth matching the channel bandwidth are transmitted. Each symbol is associated with a sequence of pulses, called burst: different symbol rates can be obtained by varying the number of pulses in a burst, while keeping constant the pulse emission rate. Data bits are mapped on symbols by using a combination of Pulse Position Modulation (PPM) and Pulse Amplitude Modulation (PAM) to transmit data. The PPM modulation is mandatory for all devices compliant to the standard, while the PAM modulation is optional. Depending on the number of pulses in a burst, on the pulse repetition rate, and on the adopted modulation scheme (PPM vs. PPM and PAM), the UWB phy can achieve bit rates varying approximately between 0.1 Mb/s and 26 Mb/s.

As mentioned before, the standard also introduces an alternative solution based on the adoption of chirp signals. The chirp solution has the advantage of working in the 2.4 GHz ISM band, thus allowing the deployment of 802.15.4a networks almost worldwide, including countries where UWB emissions are not yet allowed. The standard defines for chirp signals 14 channels spaced of 5 MHz, in the frequency range between 2410 and 2486 MHz; as a consequence of the lower bandwidth compared to the UWB case, the bit rate achievable with chirp signals is at most 1 Mb/s. Furthermore the chirp approach does not support ranging, although accurate ranging solutions based on chirp signals were proposed [5].

III. NETWORK DEVICES AND TOPOLOGIES

The 802.15.4/4a standards define two classes of devices: Full-Function Devices (FFD), in which all network functionalities are implemented, and Reduced-Function Devices (RFD), that only support a reduced set of functionalities and are thus typically sensor nodes that measure a physical parameter (temperature, light) and are capable of executing simple commands.

RFD and FFD devices organize themselves in Personal Area Networks (PANs). A PAN is controlled by a PAN coordinator, that is a device in charge of setting up and maintaining the PAN. The role of PAN coordinator can only be taken by a FFD device, while RFD devices can only join an existing PAN by communicating with the PAN coordinator. A PAN can adopt either of the two following network topologies:

- *star topology* - In this topology, devices can only exchange information with the PAN coordinator; since all communications involve the coordinator, this topology is better suited for network architectures where a device is connected to the power network, and can thus take the role of coordinator for a long time without drowning its battery power. An example of a star topology is presented in Figure 1.
- *peer-to-peer topology* - In this topology, FFD devices can communicate directly as long as they are within physical reach, while RFD devices, due to their limitations, can only connect with the PAN coordinator. An example of a peer-to-peer topology is presented in Figure 2.

The peer-to-peer topology, thanks to its higher flexibility, potentially allows the formation of more complex topologies, for example based on multiple clusters; algorithms for the creation and management of such larger network topologies are however not part of the standards.

IV. ACCESS STRATEGIES

Medium access is one of the aspects where the two standards present relevant differences. We first describe the 802.15.4 case, and we then address the modifications introduced in 802.15.4a.

Access to the medium within a piconet is based on a combination of random access and scheduled access. In both cases medium access within a PAN is controlled by the PAN coordinator that may choose between two different modalities:

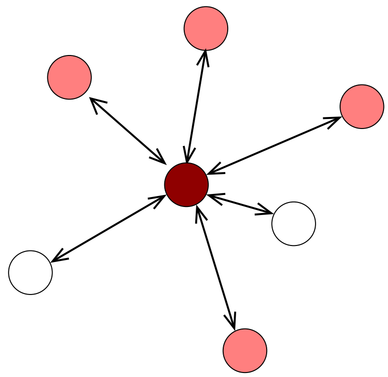

Fig. 1. Example of star topology (Dark grey circle: PAN coordinator; Light grey circle: FFD device; White circle: RFD device).

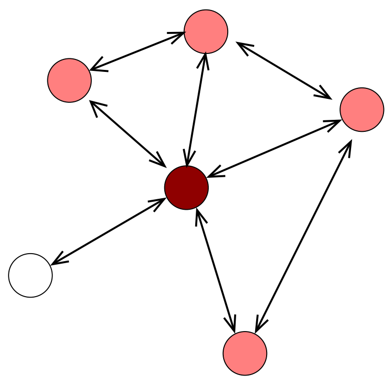

Fig. 2. Example of peer-to-peer topology (Dark grey circle: PAN coordinator; Light grey circle: FFD device; White circle: RFD device).

- *beacon-enabled*
- *nonbeacon-enabled*

In the *beacon-enabled* modality, the PAN coordinator broadcasts a periodic beacon containing information about the PAN. The period between two consecutive beacons defines a superframe structure divided in 16 slots. The first slot is always occupied by the beacon, while the other slots are used for data communication by means of random access, and form the so-called Contention Access Period (CAP). The beacon contains information related to PAN identification, synchronization, and superframe structure.

The beacon-enabled modality is adopted only when the PAN has a star topology. In this case, two data transfer modes exist:

1) *Transfer from a device to the coordinator* - a device willing to transfer data to the coordinator uses a slotted Carrier Sensing Multiple Access with Collision Avoidance (CSMA-CA). A description of the CSMA-CA protocol is provided later in this section for both slotted and unslotted versions.
   The coordinator may confirm the successful data reception with an optional acknowledgment message within the same slot.
2) *Transfer from the coordinator to a device* - when the coordinator has data pending for a device, it announces so in the beacon. The interested device selects a free slot and sends a data request to the coordinator, indicating that it is ready to receive the data. Slotted CSMA-CA is adopted to send the request. When the coordinator receives the data request message, it selects a free slot and sends data using slotted CSMA-CA as well.

In order to support low-latency applications, the PAN coordinator can reserve one or more slots that are assigned to devices running such applications without need for contention with other devices. Such slots are referred to as Guaranteed Time Slots (GTS), and they form the Contention Free Period (CFP) of the superframe. An example of superframe with both CAP and CFP is shown in Figure 3.

In the *nonbeacon-enabled* modality there is no explicit synchronization provided by the PAN coordinator. This modality is particularly suited for PANs adopting the peer-to-peer topology, but can be adopted in a star network as well.

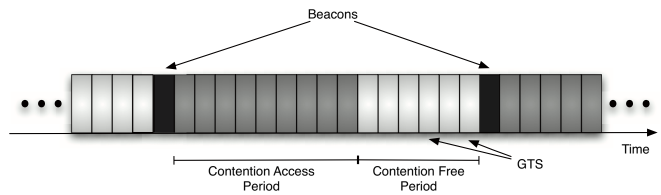

Fig. 3. Example of superframe in beacon-enabled modality.

It should be noted that the peer-to-peer topology allows for a third transfer mode: the *peer-to-peer data transfer*, in which devices exchange data without involving the PAN coordinator, thus allowing more complex topologies and larger networks.

Since there is no superframe defined in the nonbeacon-enabled modality, no GTS can be reserved, and only random access is used. Furthermore, since no slot synchronization is available, random access is adopted for medium sharing in all transfer modes. As described above, in the random access phase devices adopt a CSMA-CA protocol to access the medium, either slotted on unslotted depending on the selected PAN operation modality. It should be noted, however, that although the access protocol is referred as CSMA-CA, its implementation is actually closer to a CSMA scheme, since no Collision Avoidance packets are used. This approach differs from existing WLAN IEEE standards, such as the 802.11 family, where the access protocol foresees the use of Request To Send (RTS) and Clear To Send (CTS) in order to acquire channel control before sending data packets [6].

The CSMA-CA scheme used in the 802.15.4 standard can be summarized as follows.

In the case of a nonbeacon-enabled network, when a device needs to send data it picks a random backoff delay, defined as a multiple of a backoff time unit. When the backoff delay expires, the device performs a Clear Channel Assessment (CCA) operation, consisting in listening to the channel in order to determine if it is idle. If the channel is idle the device immediately transmits the data packet; oppositely, if the channel is busy the device repeats the procedure by picking a new backoff delay, larger than the previous one.

In a beacon enabled network the devices use a slotted version of the previous protocol to access the medium in the CAP portion of the superframe. The main differences compared to the unslotted version are the following:

1) the backoff delay unit is set to the duration of a slot in the superframe, and the beginning of the random delay interval always coincides with the beginning of a slot.
2) at the end of the random backoff delay the device performs a CCA operation at the beginning of the next slot; if the channel is idle, however, the device does not transmit the data packet immediately, but repeats the CCA for a number of slots defined by the value of a parameter called Contention Window. If the channel is idle for all the slots within the Contention Window the device transmits the data. If during one of the slots in the Contention Window the channel is detected to be busy, the device repeats the procedure by picking a new backoff delay, larger than the previous one.

The 802.15.4a standard inherits the strategy described above, with a significant difference in the channel access strategy. ALOHA was in fact introduced as an alternative channel access strategy. This decision was based on research results showing that, thanks to the Multi-User Interference robustness guaranteed by UWB, the ALOHA approach provides satisfactory throughput in UWB networks [7] for light and medium traffic loads, avoiding the additional access delay due to the collision avoidance phase. The robustness of the ALOHA approach adopted in the 802.15.4a standard is furthermore increased by the possibility of adopting a Time Hopping scheme, introducing a different delay on each burst in a packet and thus further reducing the probability of packet loss due to collisions. The CSMA-CA access was however kept as an option, in order to address high-density and high-traffic scenarios when UWB is adopted, and to enable the use of the chirp physical layer.

V. Ranging support

One of the key innovations introduced in the 802.15.4a revision is the support for ranging. The need for information on the distance between network devices was indeed one of the main reasons for the definition of the 802.15.4a standard. It should be noted however that support for ranging in 802.15.4a-compliant devices will be optional.

In order to support the broadest possible set of network topologies, the procedure to obtain distance between two devices in the network is based on a two way ranging approach, allowing for distance estimation without the need for a common time reference.

The two way ranging approach requires the two devices to exchange at least two packets, following the scheme represented in Figure 4. In this scheme, the device A starts a ranging measurement by sending a ranging packet to device B at time t_{start} and recording a timestamp. Device B records a timestamp when it receives the packet from A, and replies with a second ranging packet, transmitted after a delay ΔT when a second timestamp is recorded by B. The packet is received by the device A at time t_{stop}, when a new timestamp is recorded. The knowledge of the time interval $t_{stop} - t_{start}$ and of the delay ΔT, obtained from the four timestamps recorded by devices A and B, allows to determine the propagation time t_{flight}, given by the equation:

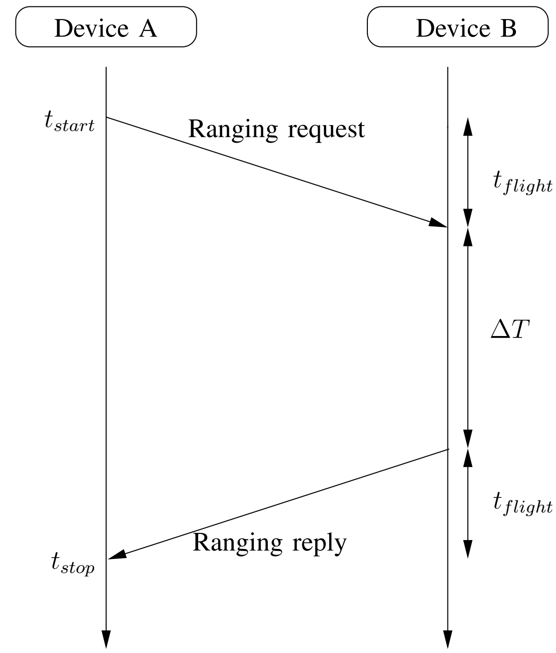

Fig. 4. Example of two way ranging. The duration of flight time was increased compared to ΔT to improve readability.

$$t_{flight} = \frac{t_{stop} - t_{start} - \Delta T}{2}. \tag{1}$$

In [8] a similar scheme was proposed for ranging in UWB networks, where the delay ΔT was assumed to be known to both devices, so that at the end of the exchange device A could determine the distance from device B. In the scheme adopted in the standard, a more general approach is proposed: the device actually performing the ranging estimation may be different from both A and B: in this case both ΔT and $t_{stop} - t_{start}$ must be communicated to the device in order to perform the ranging estimation.

Note that the two way ranging procedure involves time intervals measured by two different devices, using different reference clocks. Such clocks, although characterized by the same nominal frequency, will present different offsets from this nominal value, due to the differences in the quartz crystals used to generate the clock signals. Due to the relatively long duration of ranging packets (up to several ms) such offsets, if not taken into account in the ranging procedure, can introduce large errors in the evaluation of the propagation time. The 802.15.4a standard proposes several techniques for dealing with the offset between clocks. Depending on the implementation choices in the device, such techniques can use either specific hardware capabilities or protocol-based refinements.

On the hardware side, the measurement of the offset between the clocks in the two devices can be performed by extracting this information from the tracking loop used to keep the receiver synchronized to the incoming signal. As an example, if a device uses a digital tracking loop to maintain synchronization to the received signal, it will be forced to add or subtract units to the count in order to compensate for the offset between its own internal clock and the clock used for generating the received signal. At the end of packet reception, the number of counted units and the number of units added or subtracted to keep the lock on the received signal provide an accurate estimation of the offset between the two clocks. The capability of extracting such information and sending it to the upper layers goes under the name of crystal characterization.

Crystal characterization capability is not mandatory in the standard. If neither of the devices involved in a ranging estimation is capable of determining the offset between clocks, a protocol-based solution for compensating such offset is adopted. Such solution, based on the concept of Symmetric Double Sided Two-Way Ranging (SDS-TWR) proposed in [9], consists in repeating the packet exchange twice, inverting the role of the two devices in the second exchange. By evaluating the propagation time using the sum of the two measurements, the impact of clocks offset is highly reduced, leading to an improved ranging accuracy.

The packets used for ranging estimation are standard packets, the only difference being the value of a specific bit in the header of the packets, that identifies them as part of a ranging procedure. As a consequence, ranging can be theoretically performed without the need for any additional packet, using a data packet already scheduled for transmission, and the corresponding acknowledgment packet. In actual implementation of the standard, however, additional packets will be required in order to set-up and finish the ranging procedure, for at least two reasons:

1) enabling and disabling the ranging function in the receiving device: the recording of timestamps is performed by storing the values of a counter, leading to an additional power consumption when the ranging function is enabled. In order to minimize the impact of ranging on node autonomy, it is thus preferrable to enable this function only when a ranging procedure is required. This action requires two additional packet exchanges between devices A and B, before and after the actual ranging packets exchange.
2) preparing and concluding a private ranging procedure: the 802.15.4a standard foresees the possibility to perform a private ranging procedure, aimed at avoiding earsdropping or jamming of ranging packets by malicious devices. The private ranging procedure is based on the use of dedicated, longer preambles, that guarantee higher security for the procedure. The use of a private ranging procedure requires, however, the initiating device to communicate to the second device the use of a private preamble, and to communicate the end of the procedure after the ranging measurement is achieved.

In addition to the additional packet exchanges described above, in the case of cheap devices unable to perform crystal characterization, an additional packet exchange is required, as described above, in order to reduce the impact of clock offset and provide reliable ranging measurement.

In the worst case, ranging can thus require up to eight packets to be exchanged between the two devices involved in the procedure. Additional packets are furthermore required when the ranging procedure is requested by a third device, in order to send the ranging command to the initiating device, and to collect the time measurements from both devices involved in the procedure. The significant overhead introduced by ranging in the 802.15.4a standard should be thus taken into account in the design of applications requiring distance information, e.g. positioning algorithms, to be deployed in 802.15.4a networks.

VI. CONCLUSIONS

In this paper we provided an overview and comparison of 802.15.4 vs. 802.15.4a standards with respect to the MAC layer. We briefly described the physical layers of the two standards, highlighting the differences between the original DSSS wideband of the 802.15.4 standard and the novel UWB and chirp solutions defined in the new standard. Next, we analyzed network topology and device definitions, common to the two standards, and medium access strategies. The CSMA-CA proposed in 802.15.4 was described, and the new ALOHA approach made possible by the UWB physical layer was introduced. Finally, we analyzed in detail the ranging scheme defined in 802.15.4a made possible by the UWB physical layer, discussing the issues related to distance estimation in the standard and their potential impact on MAC procedures.

ACKNOWLEDGEMENTS

This work was partially supported by the European Union within the Integrated Project n. 506897 - PULSERS.

REFERENCES

[1] "IEEE 802.15.TG4a official web page," http://www.ieee802.org/15/pub/TG4a.html.

[2] "IEEE 802.15.4 MAC standard," Available at http://www.ieee.org/.

[3] "IEEE 802.15.4 MAC standard, 2006 revision," Available at http://www.ieee.org/.

[4] "Revision of Part 15 of the Commission's rules Regarding Ultra-Wideband Transmission Systems: First report and order," Federal Communications Commission, Tech. Rep. 02-48, April 2002.

[5] R. Lampe, R. Hach, and L. Menzer, "Chirp Spread Spectrum (CSS) PHY Presentation for 802.15.4a," IEEE P802.15 Working Group for Wireless Personal Area Networks (WPAN), Doc. IEEE P.802.15-05-0002-00-004a, 2004.

[6] "IEEE 802.11 standard," Available at http://www.ieee.org/.

[7] M. G. Di Benedetto, L. De Nardis, M. Junk, and G. Giancola, "$(UWB)^2$: Uncoordinated, Wireless, Baseborn medium access control for UWB communication networks," *Journal On Mobile Networks and Applications*, vol. 10, no. 5, pp. 663–674, October 2005.

[8] L. De Nardis and M. G. Di Benedetto, "Joint communication, ranging, and positioning in low data-rate UWB networks," in *Joint 2nd Workshop on Positioning, Navigation and Communication 2005 & 1st Ultra-Wideband Expert Talk*, March 2005, pp. 191–200.

[9] R. Hach, "Symmetric Double Sided Two-Way Ranging," IEEE P802.15 Working Group for Wireless Personal Area Networks (WPAN), Doc. IEEE P.802.15-05-0334-00-004a, 2005.