

MAC specifications for a WPAN allowing both energy saving and guaranteed delay

Part A: MaCARI: a synchronized tree-based MAC protocol

G erard Chalhoub, Alexandre Guitton, Michel Misson
{chalhoub,guitton,mission}@sancy.univ-bpclermont.fr

LIMOS-CNRS, Clermont University
Complexe scientifique des C ezeaux,
63177 Aubi ere cedex, France

Abstract. Industrials have been increasingly interested in sensor and actuator networks to monitor and control installations. The recent IEEE 802.15.4 standard has been developed to address vital issues of these networks, such as limited battery power and low processing capabilities. However, the standard does not meet all the requirements of industrial networks. For example, only some of the IEEE 802.15.4 nodes save energy, and the delay for the computer running the monitoring application to retrieve the sensor data or to activate an actuator is not bounded. Our research on energy-efficient MAC protocol is divided into two parts: Part A is the proposal of a flexible, synchronized tree-based MAC protocol called MaCARI and Part B deals with optimizations that can be performed within each cell.

This paper focuses on Part A, that is, on the description of the MaCARI protocol. MaCARI is designed to tolerate scheduled activities such as sensor data retrieval and unscheduled activities such as complex routing. MaCARI achieves this flexibility by using a tree-based centralized mechanism. We show the benefits of MaCARI by ensuring all nodes sleep regularly and by proving that the maximum end-to-end delay is bounded.

Keywords: wireless sensor networks, IEEE 802.15.4, tree-based synchronization, energy efficient MAC.

1 Introduction

With the advances in electronics, it is possible to build small, cheap, battery-powered devices that can perform basic computations, sense the environment and communicate in a wireless manner. Ideally, these devices could be deployed at a low cost and organize themselves to form a network that monitors an area of interest.

Recently, many research groups and industrials have focused on such wireless networks of sensors and actuators. Since these devices are often battery

powered, it is a vital issue to reduce the energy consumption of all the network elements (see [1] for a comparison on energy-efficient MAC protocols). The IEEE 802.15.4 standard [2] has been developed to address this problem, and has been implemented on real sensors and actuators.

The OCARI project [3] is a joint project with industrial and academic partners¹, which goal is to develop and study protocols that can increase the lifetime of a sensor and actuator network. The main scenario considered in the project is the monitoring of an industrial environment such as a factory or a production site. Such a network has the following characteristics: (i) it consists of no more than 200 sensors and actuators, with low mobility (although most of the nodes of the network are static, the propagation conditions in the environment constantly change); (ii) every network element has a limited battery-power; (iii) communications between sensors (e.g., a temperature sensor), the monitoring computer (which often is the decision maker) and actuators (e.g., an alarm bell) have a bounded delay. In the context of the OCARI project, we developed a protocol called MaCARI (MAC protocol for Ad-hoc Industrial Networks).

In this paper, we describe the MaCARI protocol. In a nutshell, MaCARI divides the time into three periods forming a global cycle: (i) the *synchronization period* allows all the network elements to be synchronized; (ii) the *scheduled activities period* is dedicated to retrieving the values of the sensors and relaying commands to the actuators; (iii) the *unscheduled activities period* can be used for running sophisticated routing protocols or simply sleeping. These three periods are shown on Fig. 1. In order to obtain the scheduling, we use a centralized approach: a specific node is in charge of creating a tree that spans all the nodes of the network. Synchronization beacons are periodically broadcasted along this tree.




Fig. 1. MaCARI divides the time into three periods: a synchronization period between T_0 and T_1 ; a scheduled activities period from T_1 to T_2 ; an unscheduled activities period between T_2 and the next T_0 . The first two periods constitute the tree-based activities. The unscheduled activities are non tree-based activities and can be used for any routing protocol.

¹ The OCARI project is a partnership between EDF R&D, DCNS, INRIA, LRI, LIMOS, One RF Technology and LATTIS with the support of the ANR (French National Research Agency).

Our contributions are three-fold. First, we merge the two solutions proposed in the 802.15.4 standard in order to avoid direct and indirect beacon collisions. Second, we reduce the energy consumption of the nodes by allowing all devices to sleep. Third, our mechanism implements multi-hop communications with delay guarantees.

This paper is organized as follows: first, we describe the 802.15.4 standard in Sect. 2. Then, we present our protocol MaCARI in Sect. 3. We provide a detailed description of the three periods, emphasizing more on the first two periods. More details on the scheduled activities component are given in Part B of this paper (please refer to [4]). After discussing some open issues in Sect. 4, we conclude our work.

2 IEEE 802.15.4 standard

The IEEE 802.15.4 standard considers two types of devices: end-devices and coordinators that are in charge of the end-devices. The MAC protocol supported in the 802.15.4 standard has two operational modes: the *non beacon-enabled* mode in which beacons are sent on request and no synchronization is required between the coordinators and their children, and the *beacon-enabled* mode in which each coordinator sends periodic beacon to synchronize the activity of its children. Only the beacon-enabled mode allows energy saving, and is therefore suited to our objectives.

2.1 Beacon-enabled mode

In this paper we are interested in the beacon-enabled mode in which the activity of the devices follows the superframe structure shown on Fig. 2. Each superframe starts with a beacon. There is an optional inactivity period between two superframes. The duration of superframes and inactivity periods is specified with the *SO* (for superframe order) and *BO* (for beacon order) parameters ($0 \leq SO \leq BO \leq 14$) contained in the beacons:

$$\begin{cases} SD = aBaseSuperframeDuration.2^{SO}, \\ BI = aBaseSuperframeDuration.2^{BO}, \end{cases}$$

where *SD* is the superframe duration, *BI* is the beacon interval, and the minimum duration of a superframe is called *aBaseSuperframeDuration*, which is equal to 15.36 ms.

During the superframe the children of a coordinator compete using slotted CSMA/CA, defined in the 802.15.4 standard [2]. An optional mechanism that offers Guaranteed Time Slots (GTS) is proposed by the standard but, to our knowledge, it has not been implemented yet due to its complexity. GTS allow end-devices to ask their coordinator for a time slot during which no other devices can be active in the cluster. All the guaranteed time slots form the contention-free period.

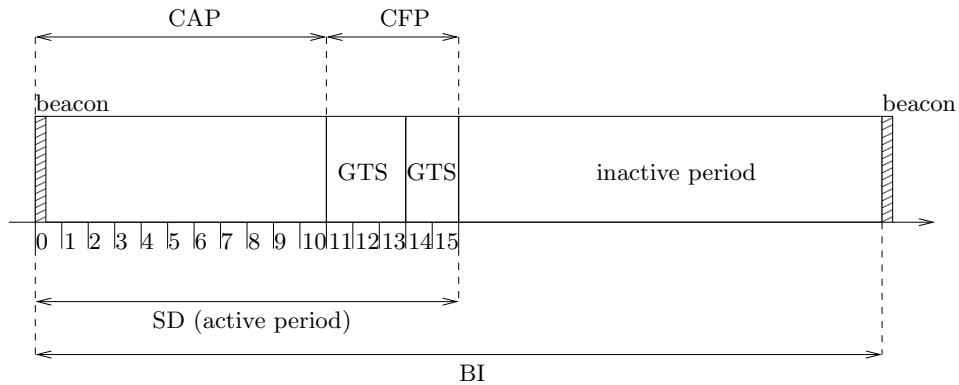


Fig. 2. The 802.15.4 superframe starts with a beacon, is followed by a Contention Access Period (CAP) during which the slotted CSMA/CA is used, and might contain a Contention-Free Period (CFP) based on Guaranteed Time Slots (GTS) allocation. If $BO > SO$, there is an inactive period that comes before the next superframe.

2.2 Cluster tree

The 802.15.4 standard has proposed to interconnect coordinators using a *cluster tree* network, as shown on Fig. 3. A cluster is formed by a coordinator and all its children (end-devices and other coordinators). The cluster tree is the union of all the clusters. A PAN (Personal Area Network) coordinator is in charge of allocating the addresses of the nodes on the tree. Details on how the cluster tree is created can be found in the 802.15.4 standard.

However, the cluster tree network suffers from beacon frame collisions when used in beacon-enabled mode. Two types of beacon frame collisions have been identified by the Task Group 15.4b [5]: *direct* and *indirect* beacon frame collisions. In both cases, two coordinators or more send their beacon frames at approximately the same time. In the first case they are in the same transmission range of each other, while in the second case they cannot hear each other but have overlapping coverage zones. The Task Group 15.4b has proposed two solutions to avoid the direct beacon frame collisions, namely the time division approach and the beacon-only period approach, and two solutions to avoid the indirect beacon frames collisions, namely the reactive approach and the proactive approach.

Solutions to direct beacon frame collisions. In the time division approach, each coordinator has to choose a time interval in a major cycle to schedule its own superframe that does not interfere with the superframes of its neighbors. An implementation for this solution was proposed in [6]. In the beacon-only period approach, all beacons are sent during a beacon-only period which precedes the superframes, each coordinator chooses an empty slot to send its beacon so that no beacon collisions occur.




Fig. 3. A cluster-tree is built on a small topology, with C_1 as a PAN coordinator. C_i represents a coordinator and E_i^x represents an end-device associated with coordinator C_i . Every coordinator constitutes a cluster with its children devices. For example, the cluster of coordinator C_2 is formed by $\{C_2, C_4, E_2^a, E_2^b\}$ and the cluster of coordinator C_3 is formed by $\{C_3, E_3^a, E_3^b\}$.

Critics have been made on these two approaches (see [7]):

- In the time division approach: (i) the duty cycles of each cluster are constrained and depend on the number of interfering clusters; (ii) communications between adjacent coordinators is not possible since their duty cycles are separated in time.
- In the beacon-only period approach: (i) the GTS mechanism is no longer possible since all the superframes are scheduled during the same time interval; (ii) dimensioning the duration of the beacon-only period is complex since it depends on the cluster tree, on the number of coordinators and on the allocation of the beacon transmission slots.

Solutions to indirect beacon frame collisions. In the reactive approach, an end-device experiencing beacon frames collisions notify its parent coordinator. This notification is received by other coordinators in range, which change their beacon transmission scheduling accordingly. In the proactive approach, the goal is to avoid indirect beacon frames collisions before they happen. This is achieved during the association procedure, that is, each time a new device joins the network: all the devices in range notify the new device of the beacon transmission time of their parents, in addition to their own beacon sending time for coordinators.

3 MaCARI protocol

In this section, we present in details the MaCARI protocol. As explained in Sect. 1, MaCARI divides the time into three periods that form a global cycle. The three periods are shown on Fig. 1.

Similarly to the 802.15.4 standard, we consider in MaCARI three types of devices: the end-devices, the coordinators and a PAN coordinator. In addition to its original role in the 802.15.4 standard, the PAN coordinator also performs global synchronization, as explained in Sect. 3.1.

MaCARI is a tree-based protocol. It uses the cluster-tree proposed by the 802.15.4 standard. An example of a small network illustrating each cluster is depicted on Fig. 3.

3.1 Synchronization period

The goal of the synchronization period, between T_0 and T_1 , is to define the global cycle by providing the same vision of a global time to all the coordinators and end-devices of the tree. This synchronization allows all devices, including the coordinators, to sleep and to wake up at predefined instants, sparing energy while sleeping.

The main difficulty is to broadcast the synchronization in a multi-hop fashion, which increases the error margins on time precision. A beacon is initiated by the PAN coordinator and propagated along the tree by the other coordinators, until it reaches all the devices of the tree.

To make sure that no collisions occur between beacons, the beacon transmission time slot of each coordinator is predefined by the PAN coordinator and included in the beacon itself. Figure 4 shows how beacons are propagated during the synchronization period for the topology shown in Fig. 3.

By T_1 , all devices should share the same global time and have their internal clocks synchronized. However, many sources of error affect this synchronization mechanism and have to be taken into consideration [8].

- the processing time before sending the beacon varies with low-level interruptions that schedule the microprocessor activities,
- the processing time after receiving the beacon varies for the same previous reason,
- the propagation delay is dependent on the distances between the devices,
- the clock drift depends on the crystal of each device internal clock.

Three solutions have been proposed by the Task Group 15.4b to reduce the error induced by these factors [5]. These solutions are based on estimating the maximum duration that each source of error might be. They could be easily implemented and included to our protocol, knowing that the distance separating the devices does not exceed 15 meters in our context and therefore the propagation delay can be neglected. The technical details concerning this issue are out of the scope of this paper.




Fig. 4. The synchronization period is a successive transmission of beacons. The PAN coordinator C_1 initiates the beacon propagation with the beacon transmission schedule (C_1, C_2, C_3, C_4) . E_1^a, E_1^b, C_2 and C_3 receive the beacon. According to the content of the beacon, C_2 is the next coordinator to propagate the beacon. C_2 sends the beacon, which is received by E_2^a, E_2^b and C_4 . Then, C_3 and C_4 propagate the beacon. By T_1 , all the devices have received the beacon. The decision to turn to sleeping mode depends on the duration of the waiting time before T_1 or the scheduled beacon transmission. For example, C_1 could decide to sleep until T_1 , while C_3 might not have time to sleep before sending its beacon.

3.2 Scheduled activities period

MaCARI schedules the activities of the devices into several activity periods between T_1 and T_2 . Each activity period concerns a coordinator and its end-devices, which form a *star*. The star is different from the cluster in that a star contains only one coordinator.

The scheduling of the activity periods of the stars provides a specific activity period to each star without interferences with the other stars (see on Fig. 5). During this period, the coordinator communicates with the end-devices. To allow the coordinators to communicate with one another, the parent coordinator is listening for the entire duration of the activity period of its children coordinators. This creates common active periods between coordinators at the end of the activity period of each child coordinator. The communications during the common active periods are depicted with arrows on Fig. 5. When these two coordinators communicate, all the end-devices of the child star are inactive.

Thus, the activity period of each star is composed of two parts: a first part during which the coordinator collects the data from the sensors or pilots its actuators, and a second part during which it can exchange data with its parent coordinator.

Each coordinator manages the activity period of its star according to the number of its end-devices and their levels of activity. The optional use of GTS is ensured without the risk of collisions caused by communications from other stars.

More details on the intra-star activity is given in Part B of this paper [4], including real measurements on Freescale components [9]. Note that we considered here that all the stars of the network were working during different time intervals; simultaneous activity periods of stars is still possible, but out of the scope of this paper.




Fig. 5. The scheduled activities period starts at T_1 and ends at T_2 . According to the content of the beacon, star of C_4 is the first to be active. At the same time, coordinator C_2 is listening and waiting for C_4 to finish communicating with its children and to initiate a parent-child communication. The same procedure applies to the other stars.

The algorithm used by the PAN coordinator to compute the size of each activity period depends on the traffic load of each star and on the application type. In the example shown on Fig. 5, we assumed that each star has the same traffic load, and the PAN coordinator therefore allocated the same activity period to each star. The same assumption has been made for the parent-child communications. Note that only the traffic with the highest priority uses these parent-child communications. The remaining traffic is forwarded during the unscheduled activities period. As for the application type, we considered that the network consists of more sensors than actuators. Thus, the activities are scheduled from the stars at the bottom of the tree to the stars at the top of the tree (refer to Sect. 4.1 to see how this algorithm can be used to guarantee an end-to-end delay).

3.3 Unscheduled activities period

The period of unscheduled activities, between T_2 and the next T_0 , is designed to allow the use of energy-efficient routing protocols. During this period, all the end-devices are sleeping (they have already exchanged data with their coordinators during the period of scheduled activities). MaCARI does not specify the activity of the coordinators: they can either be asleep or active, according to the topology

control algorithm [10] used by the routing protocol (refer to [11] for a survey on routing protocols for wireless sensor networks).

The period of unscheduled activities can also be used as a contention access period, where all the coordinators compete for the medium using a stochastic mechanism such as CSMA/CA. Such a mechanism inherently uses simultaneous transmissions, resulting in a good utilization of the channel bandwidth. However, the access to the channel is probabilistic.

The advantage of having scheduled activities in one period and unscheduled activities in another becomes apparent: messages requiring a bounded end-to-end delay are relayed during the period of scheduled activities, according to a path that could be non optimal; other messages can be routed during the period of unscheduled activities using a potentially better path.

3.4 Advantages over the 802.15.4 standard

As explained in Sect. 2, the 802.15.4 standard suffers from direct and indirect beacon frame collisions. In the context of our network specifications, MaCARI solves these problems by:

- merging the two solutions that avoid direct beacon frame collisions,
- avoiding completely the indirect beacon frame collisions.

Unlike the beacon-only period solution:

- The synchronization period (the equivalent in MaCARI of the beacon-only period) is well defined by the PAN coordinator, which knows all the coordinators.
- Since the activity periods of the stars are separated in time, the use of GTS is possible.

Unlike the time division solution, MaCARI allows communications between neighbor coordinators. They can be either parent-child communications happening in the scheduled activities period between T_1 and T_2 , or between any pair of coordinators in range during the period of unscheduled activities, between T_2 and T_0 . MaCARI does not solve the limitation on the number of clusters of the time division solution. However, this number is limited in the industrial network we focus on. Subsection 4.2 proposes a way to further reduce the impact of this limitation.

In MaCARI, there is no indirect beacon frame collision, since all beacon frame transmissions are predefined and no random choice is made for the sending time of the beacon frames.

In addition, MaCARI allows the coordinators to save their energy by sleeping during certain time intervals. This is not supported by the IEEE 802.15.4 standard.

4 Discussion

In this section, we describe a feature of the MaCARI protocol, namely, how it can be used to guarantee a bounded end-to-end delay. We also discuss how simultaneous transmissions could be used to optimize the synchronization and the scheduled activities periods.

4.1 Guaranteed end-to-end delay

In this part, we explain how high priority data can be relayed from a coordinator to another during the scheduled activities period. Low priority data is relayed during the unscheduled activities period. We prove that, under known traffic conditions, we can guarantee a bounded end-to-end delay. In the context of industrial networks, most sensors have a well defined data production, which is taken into account while planning the scheduled activities.

Let us consider the network topology shown on Fig. 3. Figure 5 represents the following scheduling of star activities (privileging communications from the bottom of the tree to the root of the tree): (C_4, C_3, C_2, C_1) , and let us assume that sensor E_4^a (E_j^i represents the i -th end-device of coordinator C_j) has to communicate to C_3 . During the activity of star C_4 , E_4^a relays its message to C_4 . Part B [4] of this paper specifies different optimizations for intra-star activities. C_4 can relay the message to its father C_2 which is listening. Later in the cycle, it is the turn of star C_2 to work. Towards the end of the activity period of star C_2 , C_2 can pass the message to its father C_1 , which is active. However, in order to communicate to C_3 , C_1 has to wait for a new global cycle to start. When it is the turn of star C_3 to work, C_1 can relay the message of E_4^a to its destination C_3 towards the end of the activity period of star C_3 .

From the scheduling (C_4, C_3, C_2, C_1) , it can be seen that messages can be relayed all the way up the tree to the root in one global cycle. However, each time a message has to go down one level on the tree, the coordinator has to wait for a global cycle.

The fact that messages can be relayed in one global cycle all the way up the tree is due to the scheduling, which always schedules the activity period of child stars before the activity period of the star of their parent. This is called an *upstream scheduling*. The reverse scheduling, which always schedules the activity period of the star of the parent before the activity period of the star of the child is called a *downstream scheduling*, since it allows to go all the way down the tree in one global cycle.

To achieve a bounded end-to-end delay, our idea consists of having an upstream scheduling and a downstream scheduling alternatively. Let us consider the worst-case end-to-end delay from a coordinator C_i to a coordinator C_j . Coordinator C_i has to send the message to the root, with a downstream scheduling for the current cycle. When this cycle ends, the message has reached only the father of C_i . Then, during the second cycle, the scheduling is upstream and the message reaches the root of the tree. Finally, during the third cycle, the schedul-

ing is downstream again and the message can reach C_j . It has taken three global cycles to forward the data from coordinator C_i to coordinator C_j .

This assumes that during the scheduled activities period, the duration of the communication between child and parent coordinators depends on the traffic generated by all the end-devices. This is taken into account by the PAN coordinator during the synchronization period, with the following constraint: each time interval allocated to a star depends on its traffic conditions, and on the position of the coordinator in the tree (coordinators close to the root have to relay more messages).

4.2 Optimizations using simultaneous transmissions

It is possible to reduce the duration of the scheduled activities period by considering simultaneous non-interfering star activities. Simultaneous activities have already been considered for intra-star communications (see the description of SGTS in Part B of this paper). This approach could be extended and applied to realize simultaneous star activities, taking into consideration the received power (and not the transmission range).

It is also possible to reduce the duration of the synchronization period by applying the same strategy. Also, if all the end-devices of a coordinator C_i receive the beacon of another coordinator C_j , C_i does not need to broadcast its own beacon.

5 Conclusions

In this part of the paper, we have presented a flexible, synchronized tree-based MAC protocol called MaCARI. MaCARI divides a global cycle into three time periods: the synchronization, the scheduled activities and the unscheduled activities periods. End-devices and coordinators are both allowed to sleep during these three periods, only to wake up and communicate during predetermined time intervals. Thus, all the network devices are able to save energy.

The period of scheduled activities allows sensors and actuators to communicate with their coordinator. Additionally, during the activity of a child coordinator, the father coordinator is active. This allows guaranteed communications between two adjacent coordinators on the tree, and therefore makes end-to-end communications possible for high priority traffic. By scheduling alternatively upstream and downstream communications, MaCARI is able to relay frames from one end of the network (e.g., a sensor) to another end of the network (e.g., a monitoring application) in three global cycles or less.

The period of unscheduled activities allows the communications between coordinators to be arbitrary (i.e., not respecting the tree structure).

Part A of this paper has focused on the synchronization and on the scheduling of star activities. More details on the optimizations of the communications within each star are given in Part B of this paper [4].

References

1. Halkes, G.P., Van Dam, T., Langendoen, K.G.: Comparing energy-saving MAC protocols for wireless sensor networks. *Mobile Networks and Applications* (10) (2005) 783–791
2. IEEE 802.15: Part 15.4: Wireless medium access control (MAC) and physical layer (PHY) specifications for low-rate wireless personal area networks (WPANs). Standard 802.15.4 R2006, ANSI/IEEE (2006)
3. The OCARI project: The ocari project web site <http://ocari.lri.fr>.
4. Livolant, E., Van den Bossche, A., Val, T.: MAC specifications for a WPAN allowing both energy saving and guaranteed delay - Part B: Optimisation of the intra-star exchanges for MaCARI. In: Submitted to WSN 2008
5. IEEE 802.15 Task Group 4b: TG4b contributions <http://grouper.ieee.org/groups/802/15/pub/TG4b.html>.
6. Cunha, A., Alves, M., Koubaa, A.: Implementation details of the time division beacon scheduling approach for ZigBee cluster-tree networks. Technical Report TR-070102, Polytechnic Institute of Porto (November 2007)
7. Koubaa, A., Cunha, A., Alves, M.: A time division beacon scheduling mechanism for IEEE 802.15.4/Zigbee cluster-tree wireless sensor networks. Technical Report TR-070401, Polytechnic Institute of Porto (April 2007)
8. Elson, J., Girod, L., Estrin, D.: Fine-grained network time synchronization using reference broadcasts. In: Symposium on Operating Systems Design and Implementation (OSDI). (December 2002)
9. Freescale: Zigbee / IEEE 802.15.4 Freescale solution. Technology General Information, Freescale Semiconductor (2004)
10. Rajaraman, R.: Topology control and routing in ad hoc networks: a survey. *ACM SIGACT* **33**(2) (June 2002) 60–73
11. Al-Karaki, J.N., Kamal, A.E.: Routing techniques in wireless sensor networks: a survey. *Wireless Communications* **11**(6) (2004) 6–28