The MAC Unreliability Problem in IEEE 802.15.4 Wireless Sensor Networks

Giuseppe Anastasi
Dept. of Information Engineering
University of Pisa, Italy
giuseppe.anastasi@iet.unipi.it

Marco Conti
IIT-CNR
National Research Council, Italy
marco.conti@iit.cnr.it

Mario Di Francesco
Dept. of Information Engineering
University of Pisa, Italy
mario.difrancesco@iet.unipi.it

ABSTRACT

In recent years, the number of sensor network deployments for real-life applications has rapidly increased and it is expected to expand even more in the near future. Actually, for a credible deployment in a real environment three properties need to be fulfilled, i.e., *energy efficiency*, *scalability* and *reliability*. In this paper we focus on IEEE 802.15.4 sensor networks and show that they can suffer from a serious *MAC unreliability problem*, also in an ideal environment where transmission errors never occur. This problem arises whenever power management is enabled – for improving the energy efficiency – and results in a very low delivery ratio, even when the number of nodes in the network is very low (e.g., 5). We carried out an extensive analysis, based on simulations and real measurements, to investigate the ultimate reasons of this problem. We found that it is caused by the default MAC parameter setting suggested by the 802.15.4 standard. We also found that, with a more appropriate parameter setting, it is possible to achieve the desired level of reliability (as well as a better energy efficiency). However, in some scenarios this is possible only by choosing parameter values formally not allowed by the standard.

Categories and Subject Descriptors

C.2.2 [**Computer-communication Networks**]: Network Protocols.

General Terms

Management, Performance, Reliability.

Keywords

Sensor Networks, IEEE 802.15.4, MAC Protocol, Reliability, Energy Efficiency, Scalability.

1. INTRODUCTION

Wireless Sensor Networks (WSNs) represent a very promising solution for a large amount of application scenarios. In recent years, the number of WSN deployments for real-life applications has rapidly increased and, based on recent studies [6, 16], it is expected to grow dramatically in the near future, especially in the fields of logistics, automation and control. This positive trend should also be favored by the adoption of two standards, recently released by the IEEE and the ZigBee Alliance, respectively. Specifically, the IEEE 802.15.4 standard [8] defines the physical and MAC (Medium Access Control) layers, while the ZigBee standard [25] covers the networking and application layers of the protocol stack.

MSWiM'09, October 26–29, 2009, Tenerife, Canary Islands, Spain.

There are three key requirements that need to be fulfilled for a credible deployment of WSNs in a real environment, i.e., *energy efficiency*, *scalability* and *reliability*. *Energy efficiency* is extremely important as sensor nodes are typically powered by batteries – with a limited energy budget – which cannot be replaced nor recharged, due to environmental or cost constraints [2]. Even when batteries can be replenished – e.g., by harvesting energy from the external environment [7] – efficient power management of sensor nodes is anyway required to achieve an adequate network lifetime. To this end, the 802.15.4 standard includes a power management mechanism, based on duty cycle, to minimize the activity of sensor nodes (see Section 3). *Scalability* is another important factor to be considered because the number of deployed sensor nodes may be very high, especially when large geographical areas need to be monitored. Finally, *reliability* is a key requirement, especially when WSNs are used for critical applications (e.g., in industrial and control applications [22]). Nevertheless, reliability can be easily compromised for a number of reasons. First, data messages containing sensor readings may be dropped due to collisions during the channel access or congestion phenomena, especially in dense sensor networks and/or high traffic conditions. Furthermore, the quality of wireless communication can be severely affected by multi-path fading in signal propagation and external interferences produced by other devices and machinery operating in the same frequency band of sensor nodes.

In this paper we focus on the unreliability introduced by the MAC protocol during the wireless channel access. This is the basic aspect to be considered when evaluating the reliability of a sensor network. If a large fraction of data is dropped by the MAC protocol during channel access, the sensor network will not be able to provide the desired level of reliability, even in an ideal scenario where transmission errors (due to fading and/or interferences) never occur, and sensor nodes never fail or run out of energy. Indeed, we show in this paper that, even in an ideal scenario, an 802.15.4 sensor network – with power management enabled – provides a very low reliability in terms of delivery ratio (i.e., percentage of data messages correctly delivered to the sink node). This may prevent a correct behavior of the sensing system, e.g., the timely detection of an event. We found that this behavior is caused by the 802.15.4 MAC protocol, which is not able to efficiently manage contentions for channel access, even when the number of contending sensor nodes is very limited (e.g., 5). Throughout, we will refer to this problem as the *802.15.4 MAC unreliability problem*.

To understand the ultimate reasons of this problem, we performed an extensive simulation analysis of the CSMA/CA (Carrier Sense

Multiple Access with Collision Avoidance) algorithm used by the MAC protocol to regulate the channel access. We also validated our simulation results through a set of measurements on a real testbed. We found that the poor performance is not due to the algorithm itself. Instead, it is caused by the default parameter setting suggested by the standard, which is definitely unsuited to sensor networks using power management. Actually, the 802.15.4 standard allows some flexibility in choosing MAC parameters as it defines a range of allowed values for each parameter. Our results show that, with an appropriate setting, it is possible to achieve the desired level of network reliability (i.e., a delivery ratio up to 100%), while increasing – at the same time – the energy efficiency. However, in some scenarios, this is possible only by using parameter values which are not compliant with the standard.

The rest of the paper is organized as follows. Section 2 discusses the related work. Section 3 introduces the 802.15.4 MAC protocol, while Section 4 describes the simulation setup used for our analysis. Section 5 shows the effects of the 802.15.4 MAC unreliability problem in single-hop (i.e., star) networks. The causes of this problem are investigated in Section 6. Based on the obtained results, in Section 7 we propose a simple solution to overcome the 802.15.4 MAC unreliability problem. Section 8 shows some experimental results that confirm and validate the previous simulation results. Finally, conclusions are drawn in Section 9.

2. RELATED WORK

Many previous papers refer to WSNs based on the IEEE 802.15.4 standard, which is considered the reference technology in this context. However, the suitability of the 802.15.4 MAC protocol for wireless sensor networks – especially with focus on reliability – has never been extensively analyzed. Actually, many papers are targeted to assess the performance of the 802.15.4 MAC protocol, mainly in terms of throughput and energy expenditure. However, in many cases they approach the problem analytically and, hence, introduce some assumptions to simplify the analysis. For example, [13] and [18] consider a star (i.e., single-hop) network, assume exponentially-distributed message generation times, and do not consider the case of simultaneous transmission attempts by all, or many, sensor nodes (e.g., after an inactive period). Under these conditions, they do not observe any MAC unreliability problem (our simulation results also confirm that, under these conditions, the delivery ratio is close to 100%). More realistic scenarios have been considered in [10, 11, 12, 24]. These papers investigate different aspects related to the 802.15.4 MAC performance. However, they do not find out any severe limitation in the MAC protocol behavior, especially in terms of MAC reliability.

The limited scalability of the 802.15.4 MAC protocol is pointed out in [23], where the authors analyze the performance of the standard in terms of throughput and energy consumption. They show that the performance of the 802.15.4 MAC is very poor when the number of contending nodes is high. A number of potential issues that can degrade the performance of the MAC protocol – including possible congestions caused by the simultaneous attempts of many nodes to access the wireless medium after an inactive period – are identified in [14] as well. However, both [14] and [23] address the above mentioned issues by suggesting modifications to the 802.15.4 MAC protocol, which makes their solution not compliant to the standard.

Issues related to the MAC unreliability, in terms of message drop probability, have been addressed in [17, 20, 21]. The authors of [21] consider a star network and assume that (**i**) all nodes attempt to transmit a message at the beginning of the active period and, (**ii**) the acknowledgment mechanism is disabled. They show that the message drop probability can be extremely high in this scenario, especially for large number of sensor nodes and message sizes. However, they do not consider the effects of using acknowledgements and retransmissions. In addition, they miss to investigate the ultimate reasons behind this behavior and, consequently, they do not propose any possible solution to fix, or alleviate, this problem. Both [17] and [20] also consider a star network topology and analyze the MAC protocol performance under saturated traffic conditions. They both find out that a large fraction of messages is dropped during the channel access, and the drop probability increases with the number of sensor nodes. [17] suggests using a larger exponential backoff delay to alleviate the problem. Similarly, [20] shows that using larger backoff parameter values can provide a significant decrease in the message drop probability, at the cost of a decreased throughput when the number of nodes is small. In both cases, however, the focus of the analysis is on the maximum achievable throughput and energy consumption. Therefore, the high message discard probability is not recognized as a major limitation and is only marginally addressed by the authors.

In this paper we thoroughly investigate the ultimate reasons behind the MAC unreliability problem, and propose a general solution for it. We use a realistic traffic model with acknowledged traffic. Finally, unlike [14] and [23], we do not propose any modification to the standard MAC protocol. Instead, we show that the MAC unreliability problem can be overcome by an appropriate setting of the MAC protocol parameters.

3. IEEE 802.15.4 STANDARD

IEEE 802.15.4 [8] is a standard for low-rate, low-power, and low-cost Personal Area Networks (PANs). A PAN is formed by one PAN coordinator which is in charge of managing the whole network, and, optionally, by one or more coordinators which control a subset of nodes in the network. Ordinary nodes must associate with a (PAN) coordinator in order to communicate. The supported network topologies are *star* (single-hop), *cluster-tree* and *mesh* (multi-hop).

The 802.15.4 standard defines two different channel access methods: a *beacon enabled* mode and a *non-beacon enabled* mode. The beacon enabled mode provides a power management mechanism based on duty cycle. It uses a superframe structure which is bounded by *beacons*, i.e., special synchronization frames generated periodically by coordinator nodes. The time between two consecutive beacons is called *Beacon Interval* (*BI*), and is defined through the *Beacon Order* (*BO*) parameter (BI=15.36ms*2^{BO}, with $0 \leq BO \leq 14$)[1]. Each superframe consists of an Active Period and an Inactive Period. In the Active Period nodes communicate with the coordinator they associated with, while during the inactive period they enter a low power state to save energy. The Active Period is denoted by *Superframe Duration* (*SD*) and its duration is defined by the *Superframe Order* (*SO*) parameter (SD=15.36ms*2^{SO}, with $0 \leq SO \leq BO \leq 14$). It

[1] Throughout the paper we assume that the sensor network operates in the 2.4 GHz frequency band.

can be further divided into a *Contention Access Period* (*CAP*) and a *Collision Free Period* (*CFP*). During the CAP a slotted CSMA/CA (Carrier Sense Multiple Access with Collision Avoidance) algorithm is used for channel access, while in the CFP communication occurs in a TDMA (Time Division Multiple Access) style by using a number of *Guaranteed Time Slots* (GTSs), pre-assigned to individual sensor nodes. In non-beacon enabled mode there is no superframe, and nodes are always active.

3.1 CSMA/CA Algorithm

The CSMA/CA algorithm is used for both the beacon enabled (during the CAP portion of the Active Period) and the non-beacon enabled modes. In the beacon-enabled mode a slotted scheme is used – i.e., all operations are aligned to backoff period slots (whose duration is 320μs) – while in the non-beacon enabled mode there is no such alignment. For brevity, in the following we will refer to the slotted scheme, highlighting the differences in the unslotted version, when necessary.

Upon receiving a data frame to be transmitted, the CSMA/CA algorithm performs the following steps.

1. A set of state variables is initialized, i.e., the contention window size (*CW*=2, only for the slotted version), the number of backoff stages carried out for the ongoing transmission (*NB*=0), and the backoff exponent (*BE*=*macMinBE*).
2. A random backoff time, uniformly distributed in the range [0, 320μs*2^{BE-1}], is generated to initialize a *backoff timer*. In the beacon-enabled mode, the start time of the backoff timer is aligned with the beginning of the next backoff slot. In addition, if the backoff time is larger than the residual CAP duration, the backoff timer is stopped at the end of the CAP and resumed at the beginning of the next superframe.
3. A Clear Channel Assessment (CCA) is performed to check the state of the wireless medium.
   a) If the medium is busy, the state variables are updated as follows: *NB*=*NB*+1, *BE*=min(*BE*+1,*macMaxBE*) and *CW*=2 (only for the slotted version). If the number of backoff stages has exceeded the maximum admissible value (i.e. *NB*>*macMaxCSMABackoffs*), the frame is dropped. Otherwise, the algorithm falls back to step 2.
   b) If the medium is free and the access mode is unslotted, the frame is immediately transmitted.
   c) If the medium is free and the access mode is slotted, then *CW*=*CW*-1. If *CW*=0 then the frame is transmitted[2]. Otherwise the algorithm falls back to step 3 to perform a second CCA.

The CSMA/CA algorithm supports an optional retransmission scheme based on acknowledgements. When retransmissions are enabled, the destination node must send an acknowledgement just after receiving a data frame. If the acknowledgment is not (correctly) received by the sender, a re-transmission is started unless the maximum number of retransmissions (*macMaxFrameRetries*) is reached. In this case the data frame is dropped.

[2] In beacon-enabled mode, before transmitting a frame, the CSMA/CA algorithm calculates whether it is able to complete the operation within the current CAP. If there is not enough time, the transmission is deferred to the next superframe.

4. SIMULATION ENVIRONMENT

To perform our simulation analysis we used the ns2 simulation tool [15], which includes the 802.15.4 module originally developed in [24] and the modifications in [19]. In all experiments we assumed that the 802.15.4 MAC protocol is operating on top of the 2.4 GHz physical layer, with a 250-Kbps maximum bit rate. The radio propagation model was two-way ground; the transmission range was set to 15 m (according to the settings in [24]), while the carrier sensing range was set to 30 m (according to the model in [1]). Unless otherwise specified, data messages flow from sensor nodes to the sink. We considered a *periodic reporting application* where sensor nodes sense the external environment and report data messages to the sink periodically, i.e., at each *communication period* (the communication period is mapped to the Beacon Period of the same duration when using the beacon-enabled mode). This is a very common case in monitoring applications. For the sake of comparison, in some experiments we also considered a Poisson message generation process. Unless stated otherwise, every sensor node generates one data message per communication period (on average, when using the Poisson process). The message size – corresponding to the MAC frame payload – is 100 bytes, while the MAC frame header is 7 bytes.

4.1 Network Scenario

We analyzed a star network (single-hop scenario) where the sink node is the PAN coordinator and all nodes (but the PAN coordinator) operate with a duty cycle for power management. More specifically, sensor nodes are placed in a circle centered at the sink node, 10m far from it. Due the considered radio model (the carrier sensing range is twice the transmission range), all nodes are in the carrier sensing range of each other. This excludes collisions due to the hidden node problem. The network uses the beacon-enabled mode, the sink acts as the PAN coordinator and all other devices as ordinary nodes associated with it. The duty cycle is set to 0.7% according to the typical values recommended by the ZigBee standard [25], which are in the range 0.1% - 2%. Specifically, the Beacon Interval is 125.8s (BO=13), while the Active Period is 0.9s (SO=6). We verified that such an Active Period is large enough to let every node send its data messages, so that the enforced duty cycle does not harm the message transmission process.

4.2 Performance Indices

In our analysis we considered the following three indices.

- *Delivery ratio*, defined as the ratio between the number of data messages correctly received by the sink and the total number of messages generated by all sensor nodes. This index measures the *network reliability* in the data collection process.
- *Average latency*, defined as the average time from when the message transmission is started at the source node to when the same message is correctly received by the sink. This index characterizes the *network responsiveness*.
- *Average energy per message*, defined as the average total energy consumed by each sensor node for each message successfully delivered to the sink. This index measures the *energy efficiency* of the sensor network.

The energy consumed by a sensor node was calculated by using the model presented in [4], which is based on the Chipcon

CC2420 radio transceiver [5]. Specifically, the model assumes the following radio states: *transmit*, *receive*, *idle* (the transceiver is on, but it is not transmitting nor receiving, i.e., it is monitoring the channel) and *sleep* (the transceiver is off and can be switched back on quickly). In addition, the model accounts for the energy spent due to state transitions as well. Although the standard does not explicitly state when the transceiver should be sleeping – except for the inactive portion of the superframe when the beacon-enabled mode is used – to further improve the energy efficiency we assume that the transceiver is in the sleep state during backoff times, as in [18].

For each experiment we performed 10 independent replicas, each consisting of 1000 communication periods. The results shown below are averaged over all replicas. The obtained confidence intervals are always very low and are thus omitted.

5. THE MAC UNRELIABILITY PROBLEM

We carried out our analysis by considering a star (i.e., single hop) network. Since many previous works have analyzed this scenario under the assumption that **(i)** sensor nodes are always active, and **(ii)** data messages are generated according to a Poisson process [10, 13, 18, 20], in this set of experiments we also considered, for comparison, Poisson message arrivals – with and without power management – in addition to the Periodic traffic.

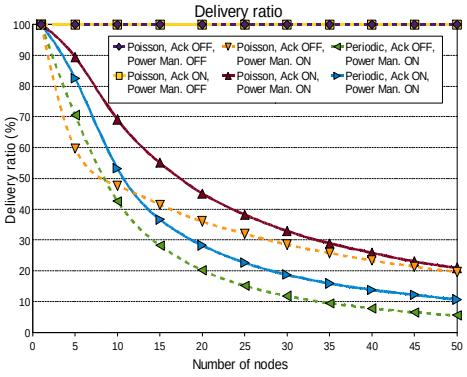

Figure 1. Delivery ratio in a star network.

Figure 1 shows the delivery ratio – as a function of the number of nodes – for Periodic and Poisson message arrivals, when the power management is enabled and disabled. When messages are generated according to a Poisson process the delivery ratio is close to 100% – even without acknowledgments and re-transmissions (curves are overlapped in Figure 1) – if sensor nodes remain always active. Instead, when power management is enabled, the delivery ratio drops sharply with the number of nodes. The delivery ratio is even lower if messages are generated periodically. As expected, retransmissions increase reliability. However, when the number of nodes is high (e.g. 50) the delivery ratio is anyway very low, i.e., around 10% and 20% with Periodic and Poisson processes, respectively.

In case of Poisson traffic, the very different behavior, with and without power management, can be explained as follows. When nodes are always active, messages are transmitted just after their generation. Since generation times are spread along the Beacon Interval there is almost no contention among sensor nodes. Instead, when power management is enabled, data transmissions are deferred to the beginning of the next Active Period, where all nodes wake up at the same time. Therefore, channel access attempts tend to become synchronized. If the message generation process is Periodic (instead of Poisson), messages are generated just before the beginning of the Active Period so as to minimize the latency. Therefore, channel accesses are perfectly synchronized and this increases the amount of contention among nodes.

We also considered different values for the Beacon Interval while leaving the Active Period constant, i.e., we varied the duty cycle. We found that the delivery ratio does not depend on the Beacon Interval when the arrival process is Periodic. Instead, it is significantly affected when messages are generated according to Poisson. In the latter case, for a given number of nodes, the delivery ratio is close to 100% when the Beacon Interval is small (i.e., the duty cycle is high) and decreases progressively as the Beacon Interval increases. The different behavior can be explained as follows. When messages are generated periodically, *all* nodes contend for channel access at the beginning of the Active Period. In case of Poisson process, since message arrivals are spread along the Beacon Interval, *not all* nodes have to contend at the beginning of the Active Period. In particular, the number of contending nodes is significantly lower than the maximum when the Beacon Interval is comparable with the Active Period.

The results in Figure 1 clearly point out that the 802.15.4 MAC protocol is not able to manage contentions efficiently, even when a limited number of nodes wake up and try to access the wireless channel simultaneously, due to power management. We also found that this problem becomes more apparent as the message size and/or message generation rate increases (see [3] for the results). Throughout, we will refer to this issue as the *802.15.4 MAC unreliability* problem.

As a final remark, it is worthwhile pointing out that the MAC unreliability problem clearly arises in cluster-tree (i.e., multi-hop) networks as well. We also performed experiments in this scenario, by using different sleep coordination schemes. Due to space limitations, the results are not presented here. The interested reader can refer to [3] for the details.

6. WHY THE PROBLEM ARISES

The results presented in the previous section show that the MAC unreliability problem may severely affect the data collection process in 802.15.4 sensor networks when power management is enabled (which occurs quite often for energy efficiency). Thus, it is very important to properly understand the ultimate reasons that originate this problem so as to remove or mitigate its negative effects. To this end, in this section we will investigate the impact of each MAC protocol parameter separately. Table 1 summarizes the MAC parameters introduced in Section 3 and the related allowed (and default) values defined in the standard (we referred to the most recent version of the standard, released in 2006).

Table 1. 802.15.4 MAC protocol parameters.

Parameter	Allowed values [8]
macMaxFrameRetries	Range: 0-7 (Default: 3)
macMaxCSMABackoffs	Range: 0-5 (Default: 4)
macMaxBE	Range: 3-8 (Default: 5)
macMinBE	Range: 0-7 (Default: 3)

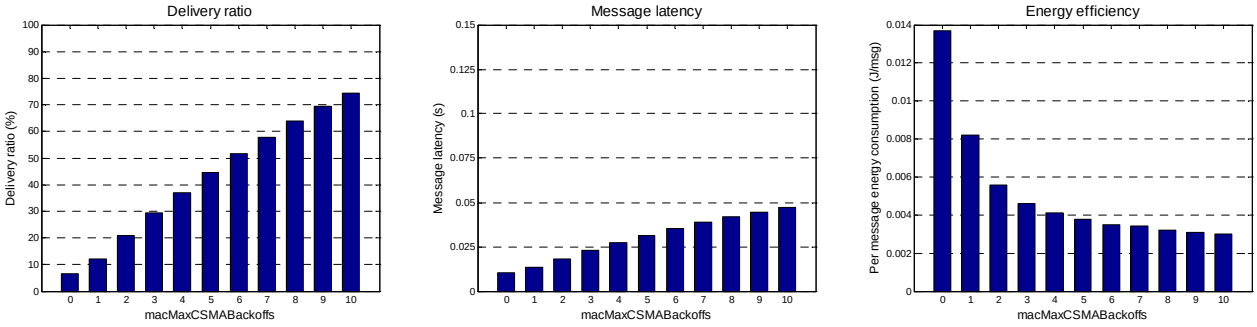

(a) (b) (c)

Figure 2. Impact of macMaxCSMABackoffs on delivery ratio (a), latency (b), and energy efficiency (c).

We focused on a star network with 15 sensor nodes and evaluated the impact of each single MAC parameter, not only in terms of delivery ratio but also in terms of energy efficiency and average message latency. The obtained results will be discussed in the next subsections. In all subsequent experiments we assumed that the message generation process is Periodic, power management is enabled and sensor nodes use acknowledgments and retransmissions to improve the reliability.

6.1 Maximum number of retransmissions

We started our analysis by investigating the impact of the maximum number of retransmissions, defined by the *macMaxFrameRetries* parameter. Since the percentage of messages successfully transmitted by sensor nodes to the sink is low (below 40% with 15 nodes), we expected that an increase in this parameter would improve the delivery ratio. Therefore, we varied *macMaxFrameRetries* in the range [0,7] (0 means that the retransmission mechanism is disabled), while setting all other MAC parameters to their default values.

Table 2. Percentage of delivered and dropped frames.

macMax FrameRetries	Delivery Ratio	Drops due to backoffs	Drops due to retransmissions
0	27.1%	59.5%	40.5%
1	33.1%	90.3%	9.7%
2	36.2%	98.2%	1.8%
3	37.1%	99.7%	0.3%
4	37.2%	100.0%	0.0%

Surprisingly, we observed that the delivery ratio is only slightly affected by this parameter (see Table 2). There is an improvement when passing from 0 to 1 or 2 retransmissions. Then, any further increase does not provide any significant effect. This happens because, in the scenario under investigation, the wireless channel is assumed to be ideal and, hence, transmission errors never occur. In addition, collisions caused by hidden nodes are not possible because all sensor nodes are in the carrier sensing range of each other (see Section 4.1). Therefore, data frames are discarded by the MAC protocol because (**i**) they collide several consecutive times with frames from other nodes and exceed the maximum number of retransmissions, or (**ii**) they exceed the maximum number of backoff stages because the CSMA/CA algorithm finds the wireless medium always busy. We measured the percentage of data frames discarded due to (**i**) and (**ii**). The obtained results are shown in Table 2. When *macMaxFrameRetries* is equal to 0 a large fraction of frames is discarded due to exceeded number of retransmissions (they are transmitted just once). However, this fraction reduces progressively and becomes negligible when *macMaxFrameRetries* is larger than 2. In any case, the majority of frames are discarded due to an exceeded number of backoff stages.

6.2 Maximum number of backoff stages

In the current and next sections we will try to understand why the maximum number of backoff stages is exceeded so frequently. To this end, it may be worthwhile to recall that the CSMA/CA algorithm performs a new backoff stage whenever the medium is found busy. At each stage the backoff window is doubled until the maximum value (defined by *macMaxBE*) is reached. Then, it remains constant. We started considering the impact of the *macMaxCSMABackoffs* parameter, which specifies the maximum allowed number of backoff stages. We varied this parameter in the range [0-10], even if the range allowed by the standard is [0-5], and set all the other MAC parameters to their default values.

The results obtained are summarized in Figure 2. As expected, the delivery ratio increases with *macMaxCSMABackoffs*. The average message latency increases accordingly as a larger percentage of messages is now delivered. Instead, the average energy per message decreases significantly when *macMaxCSMABackoffs* increases. This is because the CSMA/CA algorithm is able to successfully transmit a larger number of messages and, hence, sensor nodes use their energy more efficiently. Increasing the *macMaxCSMABackoffs* value thus results in a significant improvement in terms of network performance, but it does not solve the MAC unreliability problem. The delivery ratio remains below 80% even when 10 backoff stages are allowed. This limited contribution can be explained by observing that, when using the default values for *macMinBE* (3) and *macMaxBE* (5), the backoff window reaches its maximum value after only 3 backoff stages. Very likely, a considerable improvement could be obtained by also allowing a larger maximum backoff-window size. This intuition will be investigated in the next section.

6.3 Maximum backoff window

To analyze the impact of the maximum backoff window size we varied *macMaxBE* in [3,10] (the allowed range would be [3,8]) and set *macMinBE* and *macMaxFrameRetries* to their default values. Since *macMaxCSMABackoffs*, *macMaxBE* and *macMinBE* are bounded by the following constraint

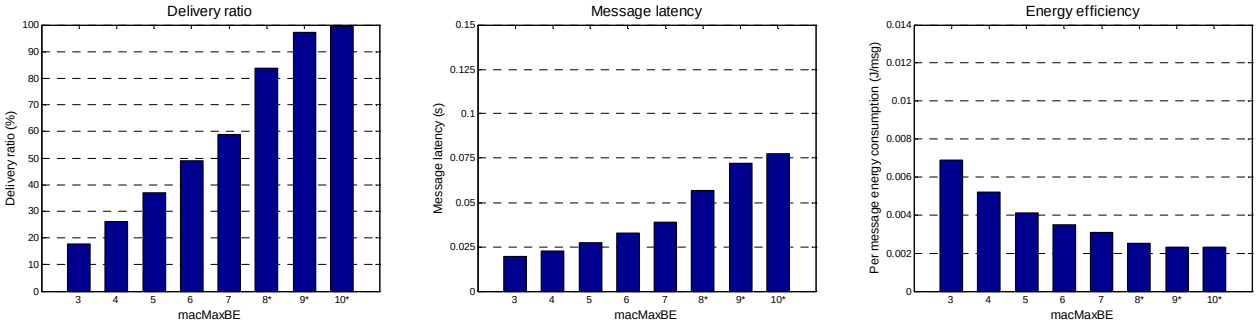

(a) (b) (c)

Figure 3. Impact of macMaxBE on delivery ratio (a), latency (b), and energy efficiency (c).

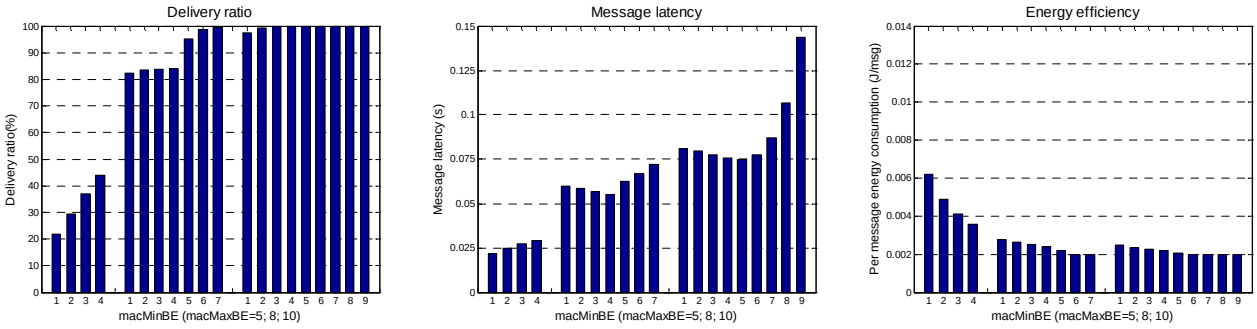

(a) (b) (c)

Figure 4. Impact of macMinBE on delivery ratio (a), latency (b), and energy efficiency (c).

$$macMaxCSMABackoffs \geq macMaxBE - macMinBE \quad (1)$$

while varying *macMaxBE* we also varied *macMaxCSMABackoffs* accordingly. The obtained results are summarized in Figure 3.

We can observe that the delivery ratio significantly improves when *macMaxBE* increases, due to the combined effect of a larger backoff window size and a higher number of backoff stages. A reliability very close to 100% is reached when using *macMaxBE*>9, while the average energy consumed per message is approximately halved when *macMaxBE* changes from 5 (default value) to 10. As expected, the average latency increases, due to the larger number of messages delivered to the sink.

6.4 Minimum backoff window

Finally, we investigated the impact of *macMinBE* which determines the minimum backoff window size. In a dense environment, where nodes wake up in a synchronized way, it might make sense to increase the value of *macMinBE* as most collisions are expected to occur in the first backoff stages. Since a correct parameter setting requires *macMinBE*<*macMaxBE*, we considered three different set of values for *macMinBE*, corresponding to three different *macMaxBE* values (i.e., 5, 8 and 10). Specifically, we fixed *macMaxBE* and varied *macMinBE* in the range [1, *macMaxBE*-1]. In addition, we set *macMaxCSMABackoffs* according to Equation (1), and *macMaxFrameRetries* to its default value. The obtained results are summarized in Figure 4. For a fixed value of *macMaxBE*, the delivery ratio and the energy efficiency tend to improve when the minimum backoff window size increases. This is because a larger initial backoff window reduces the collision probability in the first backoff stages.

From the foregoing results we can draw the conclusion that it is better to start with a very large backoff window, rather than performing many backoff stages. This is because, at the end of each backoff stage the CSMA/CA algorithm performs at least one CCA (i.e., channel assessment), which is a power-hungry operation. Instead, during the backoff time the transceiver is put in the low-power mode.

6.5 Learned lessons

Based on the above results, the following conclusions can be drawn. In the considered scenario increasing the maximum number of retransmissions does not provide additional reliability because the wireless channel is assumed to be ideal and sensor nodes are never hidden to each other. Instead, the delivery ratio can be easily improved, even up to 100%, by increasing one or more of the other MAC parameters, i.e., *macMinBE*, *macMaxBE*, and *macMaxCSMABackoffs*. The cost to be paid is an increase in message latency and total energy consumption. This is acceptable since the number of messages correctly delivered to the sink is now much higher. Moreover, the average energy consumed per correctly delivered message reduces significantly, i.e., the network become even more energy efficient. The conclusions drew above suggest that the MAC unreliability problem, observed in sensor networks where nodes tend to have a synchronized behavior (e.g., due to power management), is not intrinsic to the CSMA/CA algorithm used by the MAC protocol, but it is originated by the default MAC parameter values that have been used. It clearly emerges that the default parameter set is absolutely not appropriate for sensor networks with power management enabled. The question is, thus, whether a more appropriate MAC parameter setting can remove or, at least, alleviate the problem. This will be investigated in the next section.

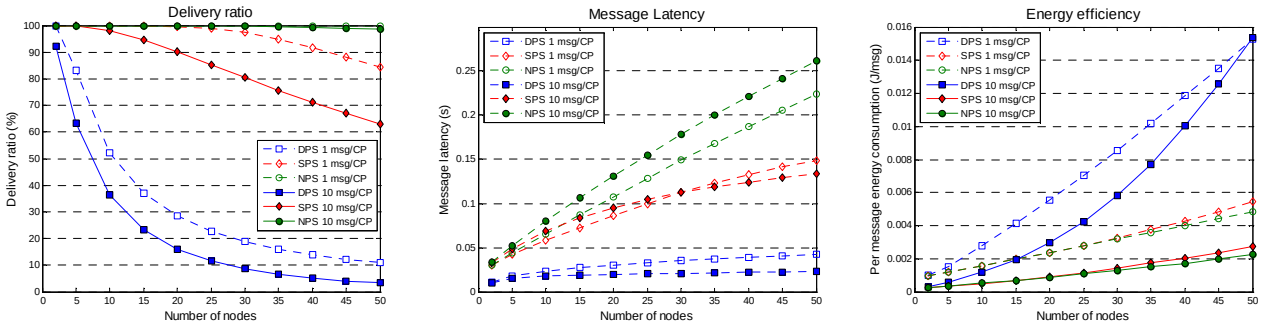

(a) (b) (c)

Figure 5. Delivery ratio (a), average message latency (b), and energy efficiency (c) in a star network, for different parameter sets.

7. POSSIBLE SOLUTIONS

To answer the previous question we considered three different sets of MAC parameter values, defined as follows and summarized in Table 3.

- *Default Parameters Set (DPS)*. This set consists of the default values specified by the standard.
- *Standard Parameters Set (SPS)*. This set consists of values still compliant with the 802.15.4 standard. In detail, *macMaxCSMABackoffs*, *macMinBE*, and *macMaxBE*, are set to the maximum values allowed by the standard, while *macMaxFrameRetries* is set to its default value.
- *Non-standard Parameters Set (NPS)*. This set of values is not compliant with the 802.15.4 standard. In particular, *macMaxCSMABackoffs*, *macMinBE*, *macMaxBE*, are set to values beyond the maximum ones allowed by the standard, while *macMaxFrameRetries* is still set to its default value.

In the following we will consider again the single-hop network already analyzed in Section 5, and will re-derive the performance under the three different parameter sets defined above.

Table 3. Different MAC parameter sets.

Parameter	DPS	SPS	NPS
macMinBE	3	7	8
macMaxBE	5	8	10
macMaxCSMABackoffs	4	5	10
macMaxFrameRetries	3	3	3

Figure 5 shows how the performance of a star network, with different number of nodes and message generation rates, changes when varying the MAC parameter set. In terms of delivery ratio (Figure 5-a), there is dramatic increase when passing from DPS to SPS. However, the delivery ratio still remains significantly below 100%, especially when the sensor network is very large and/or the workload is high. Instead, when using the non-standard parameter set (i.e., NPS) the delivery ratio is very close to 100% even in such extreme conditions. Obviously, passing from DPS to SPS (or NPS) implies an increase in the average message latency (Figure 5-b) and the total energy consumption (not shown here for the sake of space) which is, however, due to the larger number of delivered messages. Moreover, even when there are 50 nodes and the message rate is 10 messages per communication period, the average message latency is around 250 ms, which is largely acceptable for most of sensor network applications. Finally, if we consider the average energy consumption per message, instead of the total energy consumption, we can observe a significant decrease when passing from DPS to NPS or SPS (see Figure 5-c).

It is also worthwhile emphasizing that the non-standard set does not degrade the performance when the number of sensor nodes is very low (e.g., 5 or below). With respect to the default set, the additional latency introduced is below 50 ms and the energy efficiency is even improved. These results confirm that the MAC unreliability problem is caused by the default parameter values suggested by the standard and show that, in the considered scenario, a delivery ratio of 100% (or very close to 100%) can be achieved by just setting the MAC parameters to more appropriate values. However, a set of values compliant to the 802.15.4 standard may not be adequate to provide the desired level of reliability, especially when the number of nodes is large.

8. EXPERIMENTAL RESULTS

Since simulation experiments might not take into account all factors that can occur in a real environment, we also performed a set of measurements on a real sensor network testbed. The purpose of this experimental analysis is mainly the validation of simulation results presented above.

We set up an experimental testbed consisting of Jennic's 802.15.4/ZigBee JN5139 sensor nodes [9]. These nodes implement the IEEE 802.15.4 (and ZigBee) protocol stack and provide 250Kbps bit rate over the unlicensed 2.4 GHz ISM band. In our experiments we considered the same scenario described in Section 4.1. All sensor nodes use the default MAC parameters. For each experiment we performed 5 replicas. The results presented below are averaged over the 5 replicas (standard deviations are also shown).

Table 4. Comparison between the delivery ratio achieved by simulation and real experiments.

No. of nodes	4	8	12	16
Experiments	93.7 % (± 4.9%)	59.9% (± 3.4%)	42.2% (± 1.2%)	28.8% (± 2.2%)
Simulations	91.8% (± 0.4%)	61.2% (± 2.7%)	45.1% (± 0.1%)	34.8% (± 0.4%)

Table 4 compares the delivery ratio, obtained with simulation and real experiments, for different number of sensor nodes. There is a very close match between simulation and experimental results. The slightly lower values in real experiments are due to message losses (we measured a message loss in the range [0-5%] in the different experiments and replicas). We also derived other results that are not shown here for the sake of space (see [3]). Overall, they confirm and validate the simulation results discussed above.

9. CONCLUSIONS

In this paper we have investigated the performance of 802.15.4 sensor networks when power management is enabled. We have observed that, even with an ideal wireless channel, sensor nodes experience an extremely low delivery ratio. We found that the MAC unreliability problem is originated by the CSMA/CA MAC protocol, which is unable to efficiently manage contentions for channel access even when the number of sensor nodes is very limited (e.g., 5). The problem can be overcome by choosing more appropriate MAC parameter values, even though in scenarios with a large number of nodes and/or high traffic conditions, the desired level of reliability can be achieved only by using MAC parameter values out of the range allowed by the 802.15.4 standard.

We are currently extending our work to include non-ideal communication channels where transmission errors can occur (e.g., due to fading and/or interferences). Our goal is to find out the most appropriate MAC parameter set depending on the operating conditions. Since the operating conditions cannot be predicted in advance, and they also vary over time, we are also designing an adaptive scheme for adjusting dynamically the MAC parameter values, depending on the desired reliability level and actual operating conditions experienced by sensor nodes.

Acknowledgements

This work has been supported partially by the European Commission under the FP6-2005-NEST-PATH MEMORY project, and partially by the Italian Ministry for Education and Scientific Research (MIUR) under the FIRB ArtDeco and PRIN WiSe DeMon projects.

The authors wish to express their gratitude to Nicola Tucci and Luca Figoli for their help with simulation.

10. REFERENCES

[1] G. Anastasi, E. Borgia, M. Conti, E. Gregori and A. Passarella, “Understanding the Real Behavior of 802.11 and Mote Ad hoc Networks”, Pervasive and Mobile Computing, Vol. 1, N. 2, 2005.

[2] G. Anastasi, M. Conti, M. Di Francesco, A. Passarella, “Energy Conservation in Wireless Sensor Networks: a Survey”, Ad Hoc Networks, Vol. 7, N.3, May 2009.

[3] G. Anastasi, M. Conti, M. Di Francesco, “The MAC Unreliability Problem in Duty-cycled IEEE 802.15.4 Wireless Sensor Networks (Extended version)”, DII-TR-2009-04, available online at: http://info.iet.unipi.it/~anastasi/papers/DII-TR-2009-04.pdf

[4] B. Bougard, F. Catthoor, D. C. Daly, A. Chandrakasan, and W. Dehaene, “Energy Efficiency of the IEEE 802.15.4 Standard in Dense Wireless Microsensor Networks: Modeling and Improvement Perspectives”, Proc. DATE05, Volume 1, pp. 196-201, March 7-11, 2005.

[5] Chipcon CC2420 Website, http://focus.ti.com/docs/prod/folders/print/cc2420.html

[6] Embedded WiSeNTs Consortium, “Embedded WiSeNts Research Roadmap (Deliverable 3.3)”, www.embedded-wisents.org.

[7] IEEE Pervasive Computing, Special issue on “Energy Harvesting and Conservation”, Vol. 4, N.1, January 2005.

[8] IEEE Standard for Information technology, Part 15.4; Wireless Medium Access Control (MAC) and Physical Layer (PHY) Specifications for Low-Rate Wireless Personal Area Networks (LR-WPANs), IEEE Computer Society, 2006.

[9] Jennic, http://www.jennic.com/products/development_kits/

[10] A. Koubaa, M. Alves, E. Tovar, “A Comprehensive Simulation Study of Slotted CSMA/CA for IEEE 802.15.4 Wireless Sensor Networks”, Proc. WFCS’06, Torino, Italy, June 2006.

[11] K. Leibnitz, N. Wakamiya, M. Murata, “Modeling of IEEE 802.15.4 in a Cluster of Synchronized Sensor Nodes”, Proc. ITC-19, Beijing, China, August 2005.

[12] G. Lu, B. Krishnamachari, C. Raghavendra, “Performance Evaluation of the IEEE 802.15.4 MAC for Low-rate Low-power Wireless Networks”, Proc. EWCN’04, 2004.

[13] J. Mišic, S. Shafi, and V. B. Mišic, “The Impact of MAC Parameters on the Performance of 802.15.4 PAN”, Ad Hoc Networks, Vol. 3, N. 5, pp. 509–528, 2005.

[14] J. Mišic, S. Shafi, and V. B. Mišic, “Performance limitations of the MAC layer in 802.15.4 Low Rate WPAN”, Computer Communications, Volume 29, N. 13-14, August 2006.

[15] Network Simulator Ns2, http://www.isu.edu/nsnam/ns.

[16] ON World Inc, “Wireless Sensor Networks – Growing Markets, Accelerating Demands”, July 2005, http://www.onworld.com/html/wirelesssensorsrprt2.htm.

[17] S. Pollin, M. Ergen, S. Ergen, B. Bougard, L. Van der Perre, I. Moerman, A. Bahai, F. Catthoor, “Performance Analysis of Slotted Carrier Sense IEEE 802.15.4 Medium Access”, IEEE Trans. Wireless Comm., Vol. 7, N. 9, September 2008.

[18] I. Ramachandran, A. K. Das, S. Roy, “Analysis of the Contention Access Period of IEEE 802.15.4 MAC”, ACM Trans. on Sensor Networks (TOSN), Vol. 3(1), March 2007.

[19] I. Ramachandran, “Changes Made to the IEEE 802.15.4 NS-2 Implementation”, available at http://www.ee.washington.edu/research/funlab/802_15_4/ns2_changes.pdf.

[20] C. K. Singh, A. Kumar, P. M. Ameer, “Performance Evaluation of an IEEE 802.15.4 Sensor Network With a Star Topology”, Wireless Networks, Vol. 14, N. 4, August 2008.

[21] F. Shu, T. Sakurai, M. Zukerman and H. L. Vu, “Packet Loss Analysis of the IEEE 802.15.4 MAC without Acknowledgment”, IEEE Communication Letters, vol. 11, N.1, January 2007.

[22] A. Willig, “Recent and Emerging Topics in Wireless Industrial Communications: a Selection”, IEEE Transactions on Industrial Informatics, Vol. 4, N. 2, May 2008.

[23] K. Yedavalli, B. Krishnamachari, “Enhancement of the IEEE 802.15.4 MAC Protocol for Scalable Data Collection in Dense Sensor Networks”, Proc. WiOpt08, Berlin, Germany, March 31 - April 4, 2008.

[24] J. Zheng, M. J. Lee, “A Comprehensive Performance Study of IEEE 802.15.4”, IEEE Press Book, 2004.

[25] ZigBee Alliance, The ZigBee Specification version 1.0