Collision-Free Downlink Scheduling in the IEEE 802.15.4 Network

Sangki Yun and Hyogon Kim

Korea University

Abstract. IEEE 802.15.4 is the Low-Rate Wireless Personal Area Network (LR-WPAN) standard that is suitable for wireless sensor networks and wireless home networks among others. The IEEE 802.15.4 is specifically designed for energy efficiency since many 802.15.4-compliant devices are expected to operate on battery. Because of inadequate design of the downlink frame transmission mechanism in the standard, however, 802.15.4 devices can waste their energy due to collisions even under modest downlink traffic in the network. In order to solve this problem, we propose a novel mechanism which evenly distributes the downlink frame transmissions and decimate collisions. It exploits the information already imbedded in the beacon frame, so it does not require modifications of the IEEE 802.15.4 standard. Our scheme significantly reduces the energy consumption of 802.15.4 WPAN devices under modest to heavy downlink traffic, while not adversely affecting the system performance under low utilization.

Key words: 802.15.4, downlink collision, WPAN, ZigBee, energy consumption, battery

1 Introduction

The IEEE 802.15.4 is a medium-access (MAC) and physical (PHY) layer standard for low-power, low-rate wireless communication [1], which is widely considered to serve the needs of sensor networks and home networks well. Many 802.15.4 devices are expected to operate on batteries, so energy efficiency is a prime concern in the design of 802.15.4 MAC. The 802.15.4 devices can turn off the radio transceiver when they do not have frames to send, only turning it on for periodic beacon frames. If the beacon frame notifies the device of a pending (downlink) data, it issues a Data Request frame and turns on the receiver until it receives the data. But if there are multiple devices to receive such downlink frames, according to the current 802.15.4 standard, they almost concurrently attempt the transmission of request frames, leading to high collision probability that leads to additional energy consumption and delay.

In this paper, we address the problem while not modifying the 802.15.4 standard specification. Instead, our approach simply utilizes the information already embedded in the beacon frame, so 802.15.4-conformant WPAN devices can implement it without causing the interoperability issues. In our scheme, each

WPAN device is implicitly assigned its own superframe slot for the request frame transmission, which completely prevents the collisions among request frames. In contrast to the standard scheme that suffers higher collision and drop probabilities as the downlink traffic increases, our scheme has consistently low collision and drop probabilities and eventually low energy consumption regardless of the volume of the downlink traffic.

There is a dearth of prior work addressing this problem, because the IEEE 802.15.4 has been standardized only recently and this problem has not received much attention. Misic *etal.* [5] is the only one that notices the problem, but it simply suggests reducing the size of the Pending Addresses field to 3 or 4 to decrease the number of simultaneously requesting devices. Obviously, it would not only restrict the capacity and flexibility of the downlink traffic but also require the modification of the current standard. In contrast, our scheme proposed in this paper neither requires standard modification nor restricts the capacity of the downlink traffic.

The rest of the paper is organized as follows. We draw on the relevant parts of the IEEE 802.15.4 standard in Section 2. In Section 3 we characterize the downlink inefficiency problem of the IEEE 802.15.4, and introduce our solution to this problem. Section 4 presents the experimental evaluations. Section 5 concludes the paper.

2 Background

The IEEE 802.15.4 standard defines the PHY and MAC sublayer specifications for low data rate wireless connectivity with portable devices. It has data rates of 250kb/s, 40kb/s and 20kb/s. There are two types of devices in 802.15.4 WPAN networks: a full-function device (FFD) and a reduced-function device (RFD). The FFD has more computation capability and energy than the RFD and has a responsibility to be a PAN coordinator. An RFD is intended to be used in simple applications such as a light switch and a passive sensor, so it has minimal resources and memory capacity. An FFD can communicate with RFDs and other FFDs but an RFD can talk only to FFD. One FFD in a communication network is elected for a PAN coordinator. The coordinator has duties to transmit beacon frames, schedule a channel allocation, associate newly appearing devices. So usually the FFD which has consistent power supply and high computation capability (*e.g.* personal computer) becomes the coordinator. The standard specifies two network topologies: star topology and peer-to-peer topology.

2.1 Channel Structure

In a beacon-enabled network, an active period between a beacon interval is called superframe. It is divided into 16 equally sized superframe slots and organized into the Contention Access Period (CAP) and the Contention Free Period (CFP). The standard stipulates that a minimum of 9 slots should be used for the CAP. In the CAP, the device accesses the channel using slotted Carrier Sensing Multiple

Access with Collision Avoidance (CSMA/CA). In the CFP, on the other hand, the coordinator assigns the superframe slot called Guaranteed Time Slot (GTS) to each device. The length of the superframe is determined by Superframe Order (SO), and the beacon interval by Beacon Order (BO). If the value of SO and BO is equal, the superframe length and the beacon interval becomes identical, so there is no inactive period between beacons. If the BO has a higher value than the SO, the gap between the superframe duration and the beacon interval becomes an inactive period in which all devices get into the power-down mode. One superframe slot is divided into several backoff periods, and the number of backoff periods in a superframe slot depends on the SO value.

2.2 Channel Access Mechanism

The channel access mechanism in the CAP is similar to that of the IEEE 802.11 DCF, with a few notable differences. The 802.11 device always turns on the receiver (barring power saving mode), so it performs Clear Channel Assessment (CCA) in every slot. But the 802.15.4 device turns on the receiver and performs the CCA only after the end of the random backoff, in order to save energy. Because 802.15.4 devices do not perform the CCA in every slot, they cannot freeze their backoff counter even when other nodes are transmitting, so keep decrementing it. When a device finishes the backoff countdown, it performs two CCAs in a row and transmits the frame if the channel is idle. If the channel is busy, it notches up the backoff stage unless it reaches the maximum. It is a major difference with 802.11 DCF: the backoff stage increases not upon the collision, but upon the busy channel. Collisions do not influence the backoff stage. Once the devices access channel and transmit frames, they reset their Backoff Exponent (BE) value regardless of the fate of the transmission.

2.3 Data Transfer Model

The data transfer mechanism is asymmetric between the coordinator and the device. The transmission from the device to the coordinator (*i.e.*, uplink) is straightforward. If the device has a frame to send to the coordinator, it simply transmits it using the CSMA/CA. If the coordinator successfully receives the frame, it sends an acknowledgement (ACK) frame to the device.

The transmission from the coordinator to the device (*i.e.*, downlink) on the other hand uses an indirect transmission mechanism. Fig. 1 illustrates the mechanism. The coordinator notifies the devices of the pending frames through the Pending Address field in the beacon frame. The length of the field is variable, but up to 7 devices can be addressed [1]. The notified devices must transmit the Data Request frame, and the instance of the transmission is dependent on the *macAutoRequest* parameter. If it is set true, which is the default value, the device should transmit the request frame in the immediately following superframe slot, within which backoffs are performed in the units of backoff periods (BPs). Upon successfully receiving the request frame, the coordinator transmits

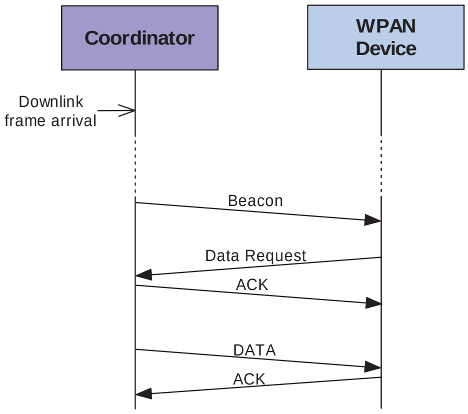

Fig. 1. Indirect transmission

the downlink frame using the CSMA/CA. After receiving the frame, the device sends an ACK frame to the coordinator.

In the indirect transmission, four frames are exchanged to transmit one data frame, so it looks inefficient. But there is a reason to use this indirect transmission: to turn off the device's radio receiver. If the downlink frame is transmitted by the direct transmission mechanism, the device must always turn on the receiver because it can not predict when the coordinator transmits the frame. On the other hand, in the indirect transmission, the device can turn on the receiver only for the beacon frame and immediately turn off if no pending frame exists.

3 Proposed Idea

3.1 The Problem of the Indirect Transmission

The indirect transmission achieves energy savings, but has a drawback. The collision probability easily reaches a high value for even a modest number of notified stations. Under the default macAutoRequest parameter setting, as soon as the beacon frame is transmitted, the involved devices jump into contention. The initial contention window is given by the macMinBE parameter, whose default value is upper bounded by 3 (it implies that initial maximum contention window size is 7). Thus if even 3 or 4 devices out of the maximum 7 contend, the collision and backoff probabilities become significant.

In addition to the high collision and backoff probabilities, the feature has one more negative impact: it elongates the duration for which the device's radio receiver remains on. As discussed earlier, after a device successfully transmits the request frame, it must turn on the receiver until the arrival of the data frame. Staying longer in the receiving mode is critical for battery life, because being in the receiving mode can consume comparable or even more energy than transmission [3]. So the 802.15.4 standard stipulates that the device can wait the downlink frame for only up to 61 backoff periods. If the downlink frame is not received within this period, the device turns off the receiver and transmits the request frame in the next superframe again.

Table I shows the energy consumption when the device stays in each state for the duration of one backoff period, based on the popular TI Chipcon 2420 chipset data sheet [3]. The data indicates that receiving mode consumes slightly more energy than the transmission mode.

Table 1. Energy consumption in each state

State	**Energy consumption (μJ/sec.)**
Power Down (Sleep mode)	0.00064
IDLE	0.136
TRANSMISSION	5.57
RECEIVING	6.02

3.2 Time-ordered Slot Appointment Rule (TSAR)

Our solution approach to the Data Request collision problem is to utilize the information about the number of pending devices that is already available in the beacon frame. Every device listens to the beacon frame so can notice how many devices will request the downlink transmission in the upcoming superframe by reading the Pending Addresses fields. By utilizing this information, we can distribute each device's contention period, and eventually lower the collision and drop probability. Below, we will call this scheme Time-ordered Slot Appointment Rule (TSAR) for convenience.

In TSAR, each WPAN device examines in the beacon frame to find where its address is positioned. If a device's address is located at n^{th} position in the Pending Address field, it starts its contention at the n^{th} superframe slot from the left boundary. As discussed earlier, one superframe has 16 superframe slots of which the minimum of 9 slots are guaranteed for the CAP, and a beacon frame can convey up to 7 Pending Addresses. So all devices can be assigned their own superframe slots. This significantly reduces the collision probability for the Data Request frames.

Fig. 2 and Fig. 3 compares the implied behavior in the standard with that of TSAR, when 3 devices are notified of the pending frame. In the figures, B means backoff duration. In the standard scheme, every device starts the contention from the first backoff period after the beacon frame. Therefore, they contend with each other and have high collision and backoff probabilities. On the other hand, in TSAR, each device starts its backoff in different superframe slot, so they are isolated in terms of contention. Although the uplink traffic can interfere, we will see later that such isolation is for the most part retained.

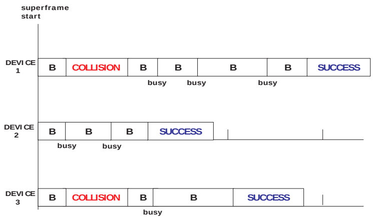

Fig. 2. Channel access in the standard (macAutoRequest)

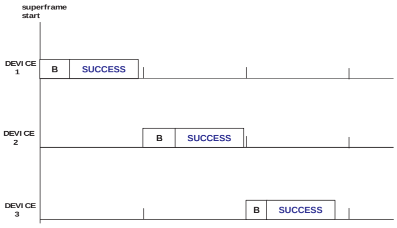

Fig. 3. Channel access in TSAR

4 Performance Evaluation

In this section, we compare TSAR with the standard scheme. Although NS-2 [4] provides an environment for simulating 802.15.4, it poses difficulties for measuring the energy consumption and the delay per downlink frame. So for the experiments below, we use a home-brewed event-based simulator. Parameters used in simulations are summarized in Table 2. We assume that the WPAN has the star topology where a single coordinator communicates with 10 devices. The Superframe Order (SO) is set to 4, and the Beacon Order (BO) is set equal to the SO unless otherwise noted. We vary the downlink packet arrival rates while fixing the uplink packet arrival rate. In our experiments, the arrival rate represents the number of arriving packets over the entire system (not per device) in a given time.

Table 2. Simulation settings

Parameters	**Values**
Simulation time	300 second
Transmission rate	250 Kbps
macMinBE	3
Topology	star-topology
Number of devices	10
Payload size	90 bytes
Superframe Order	4
Beacon Order	4, 5
Maximum backoff limit	5

4.1 Collision and Drop Probability

In this experiment, we measure two probabilities. First, P_c is the probability that the request or data frame collides with the transmission from other node(s). The second is P_d, the probability that a device gives up transmission due to repeated failures over limit.

Fig. 4 shows the P_c and P_d of the standard (labeled *macAutoRequest*) scheme and TSAR as functions of the downlink packet arrival rate, where there is no uplink traffic. The standard scheme suffers from the gradually increasing collision and blocking probability, whereas TSAR completely avoids either collisions or drops. It shows that TSAR successfully resolves channel accesses at least between downlink frames. Fig. 5 measures the probabilities in face of varying uplink traffic intensity. We set the uplink packet arrival rate R_u to 10, 30, and 50. Even when the uplink traffic interferes, TSAR maintains consistently low collision and drop probabilities compared with the standard scheme.

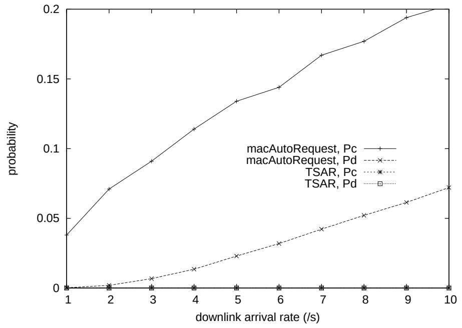

Fig. 4. Collision and drop probability, no uplink traffic.

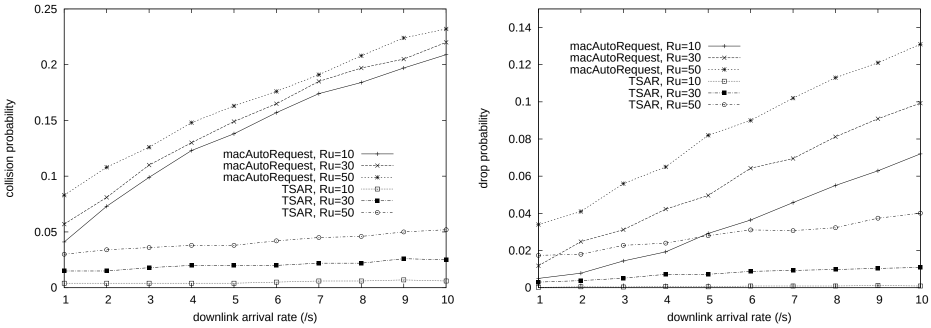

Fig. 5. Collision and drop probability with varying uplink traffic.

So far, we have set SO equal to BO. Here, we let SO be smaller than BO so there is an inactive period between superframes. Given the same packet arrival rate, it implies the traffic intensity and the contention level rises higher in this set of experiments since the traffic focuses on the active period. In particular, we set $SO = 4$ and $BO = 5$, *i.e.*, 50% duty cycle. So the superframes and the inactive periods each have 250ms duration. Fig. 6 shows the collision and drop probabilities.

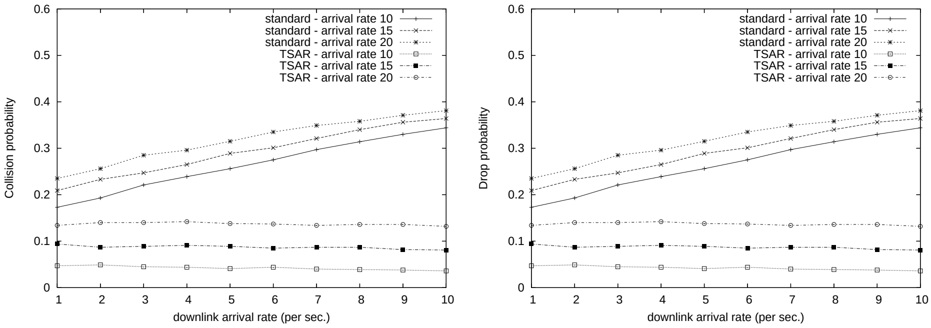

Fig. 6. Collision and drop probability with inactive period and uplink traffic.

The results are qualitatively similar to those without the inactive period, implying that the inactive period does not adversely affect TSAR. A noticeable phenomenon in the figure, however, is that TSAR has slightly decreasing collision probability, whereas the standard scheme has increasing collision probability with the downlink arrival rate. Note that with the idle period, the uplink frames arriving during the idle period accumulate, and they are launched as soon as the next superframe starts. In contrast, in the absence of the idle period they are spread over the superframe. The consequence is that with the idle period, the Data Request frames in the latter part of the superframe becomes less interfered by the uplink transmissions. Also note that the increase of the downlink traffic intensity means more and more superframe slots are utilized by TSAR, making the probability of encounter with the uplink transmissions lower as we go to the latter part of the superframe. So the average probabilities drop in Fig. 6. Although this effect can mitigate with intensified uplink traffic, it definitely contributes to the stability of the TSAR scheme.

4.2 Energy Consumption and Delay

One aspect of TSAR is that it can increase the delay for downlink data transmission as it intentionally moves all but one Data Request frame transmissions after the first superframe slot. But then again the backoff and collision probabilities are lower than the standard scheme, affecting the delay. So it is not straightforward to estimate the impact of TSAR on the delay performance. In this section, we weigh the benefit and cost of TSAR by comparing its delay and energy preservation performance with that of the standard scheme.

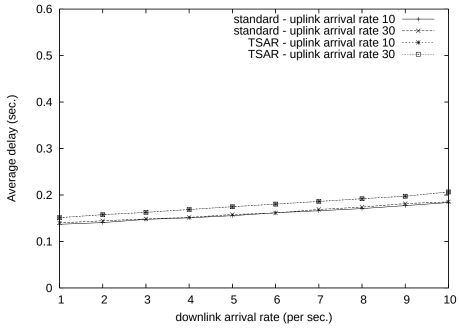

Fig. 7. Average delay comparison.

Fig. 7 compares the average downlink delay of the standard scheme with that of TSAR. Due to the indirect transmission, the downlink frame should wait for the next beacon frame for it to be notified to the device. Assuming uniform packet arrival, the mean of this delay is half the beacon interval. In this experiment, we set $BO = 4$ so it is 125ms in our setting. This delay is constant regardless of the arrival rate. After the beacon frame, the downlink frame waits for the request frame from the device and then it is transmitted. This delay is influenced by the backoff and collision probabilities and the instance that the device jumps into contention for the request frame transmission, so depends on the channel access mechanism. From the simulation result, we observe that the former delay that TSAR does not affect contributes the majority of the total delay. TSAR adds only about 10ms to what the standard scheme already incurs. This delay difference is negligible, considering the delay requirements of most IEEE 802.15.4 WPAN applications [2].

Fig. 8 addresses the energy aspect. In this experiment, we measure the energy that the device consumes to receive one downlink data frame. The energy consumption in each state is based on Table 2 [3]. We observe that the energy consumption of the standard and TSAR is similar when the downlink traffic is low, but the gap grows as it intensifies. TSAR performance is about 50% superior when downlink arrival rate is 5 frames per second compared with the standard. It is mainly because the device with the standard scheme stays longer with its receiver on than the device with TSAR. In the standard scheme, requesting devices contend with each other with increasingly larger number of collisions with traffic intensity so the time to receive their downlink frames grows. But in TSAR, since the requesting devices do not contend with each other, the probability of such event happening is much lower even considering the interference from the uplink traffic.

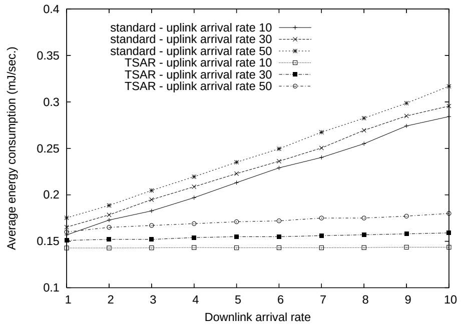

Fig. 8. Average energy consumption comparison.

Acknowledgement

This work was supported by grant No. R01-2006-000-10510-0 from the Basic Research Program of the Korea Science & Engineering Foundation.

5 Conclusion

The IEEE 802.15.4 is a WPAN PHY/MAC standard that regards the efficient use of the battery as a prime consideration. But the indirection transmission mechanism in the 802.15.4 turns out to harbor inadequacy that causes energy waste when moderate downlink traffic exists such as 3 or 4 frames per superframe. Our scheme, TSAR, solves this problem without the modification of the standard specifications. As such, our scheme can be implemented in the ZigBee-compliant devices. The experiment results show that TSAR always saves energy over the current standard with minimal delay cost. TSAR achieves this without adversely affecting the performance when the traffic is light.

References

1. IEEE Standard for Part 15.4: *Wireless Medium Access Control Layer (MAC) and Physical Layer (PHY) Specifications for Low Rate Wireless Personal Area Networks (LR-WPANS)*, Oct. 2003
2. ZigBee Alliance, http://www.zigbee.org.
3. Chipcon data sheet for 2.4GHz 802.15.4. Available at http://www.chipcon.com/files/CC2420_Data_Sheet_1_3.pdf.
4. The ns-2 simulator. Available at http://www.isi.edu/nsnam/ns/.
5. Jelena Misic., Shairmina Shafi, and Vojislav B. Misic, "Avoiding the Bottlenecks in the MAC Layer in 802.15.4 Low Rate WPAN", in Proceedings of IEEE ICPADS 2005.