Computer Communications 81 (2016) 86–96

Contents lists available at ScienceDirect

Computer Communications

journal homepage: www.elsevier.com/locate/comcom

Mutual interference in large populations of co-located IEEE 802.15.4 body sensor networks—A sensitivity analysis

Amirhossein Moravejosharieh, Andreas Willig*

Department of Computer Science and Software Engineering, University of Canterbury, Private Bag 4800, Christchurch 8140, New Zealand

ARTICLE INFO

Article history:
Received 13 May 2015
Revised 1 January 2016
Accepted 5 January 2016
Available online 12 January 2016

Keywords:
Wireless body sensor networks
IEEE 802.15.4
Internal interference
Sensitivity analysis

ABSTRACT

We consider scenarios where a large number of wireless body sensor networks (WBSN) meets at the same location, as can happen for example at sports events, and assess the impact of their mutual interference on their achievable transmission reliability. In particular, we consider several of MAC- and application parameters for a range of static and dynamic schemes for allocating WBSNs to frequencies, and determine their relative impacts on achievable performance. Our results indicate that parameters related to the MAC backoff scheme have by far the largest impact on performance, and that frequency adaptation can provide substantial performance benefits.

1. Introduction

Wireless Body Sensor Networks (WBSNs) are expected to play a pivotal role in health-related and well-being applications [1–5]. They are deployed to measure and collect human vital signs for diagnosis and monitoring of medical conditions or assessment of training progress. Key characteristics of WBSNs are their relatively small size (both in number of sensors and the network diameter), mobility of a WBSN as a whole, and the often strict requirements in terms of reliability and timeliness for transmission of vital data.

The IEEE 802.15.4 standard [6,7] is a well-established standard for low-power wireless sensor networks, which has also been considered as underlying technology for WBSNs, not the least due to the availability of cheap and mature components. The IEEE 802.15.4 standard covers both the medium access control (MAC) and physical layers (PHY). On the physical layer the standard supports different frequency bands, with the 2.4 GHz ISM band being arguably the most popular one. The IEEE 802.15.4 standard partitions this band into 16 frequency channels and the standard suggests that a WBSN picks one of these channels and stays there. In this frequency band a WBSN can be subjected to external interference coming from other technologies like for example WiFi or Bluetooth, and this can impact the achievable reliability and timeliness considerably [8–10].

In this paper we consider another type of interference which is fundamentally different from external interference (which is often considered to be equivalent to noise), and this is *internal interference*, i.e. interference coming from co-located networks of the same technology and sharing the same frequency band. One of the key differences between internal and external interference is that normally very little information can be extracted out of external interference, it is essentially the same as noise. In contrast, an IEEE 802.15.4 WBSN can collect quite useful information from internal interference: it becomes possible to receive packets (in particular beacons) and gather information about the number of other WBSNs on the same channel, their beacon periods, and so forth. This information can be used to adapt physical layer, MAC layer or application parameters.

We address situations where many people wearing WBSNs gather at the same place, for example a sports event, in a cafe, a concert or theater performance and others. All these application scenarios have in common that they lead to a very high density of WBSNs. The WBSNs of different people are completely unsynchronized and will compete with each other to gain access to the frequency spectrum and time resources, and there is a risk that many of them will not be able to achieve the desired reliability and timeliness. Following up on previous work [8], we hypothesize that giving a WBSN the ability to adapt its frequency channel over time might be very helpful to deal with internal interference.

In this context we consider a few important questions. The first main question addressed in this paper is: How many WBSNs can meet at the same place so that only a small percentage of them experiences un-acceptable performance degradations in terms of packet loss rates? We will define precise performance measures capturing this question and which we will refer to as the *satisfaction rate* and the *carrying capacity*. We will evaluate these

---

* Corresponding author. Tel.: +64336429877869.
E-mail address: andreas.willig@canterbury.ac.nz (A. Willig).

http://dx.doi.org/10.1016/j.comcom.2016.01.002
0140-3664/© 2016 Elsevier B.V. All rights reserved.

performance measures by simulation for a range of schemes (some of which we have proposed in our previous work, see [8,11]), which either choose their operating frequency only once during initialization, or which can adapt their operating frequency dynamically. Our results suggest that schemes with the ability to dynamically adapt their operating frequency and making careful choices about their next frequency can provide substantial improvements over schemes which do not adapt their frequency.

The second main question is: How sensitive are satisfaction rate and carrying capacity against variations of several important system parameters like for example the traffic load, the beacon generation period of the WBSNs, or the MAC backoff parameters? To answer this, we apply the response surface methodology (sometimes also referred to as a 2^k full factorial experiment, see [12,13]) for the satisfaction rate and identify the parameters contributing most to observed variation. Our results indicate that in particular the parameters of the MAC backoff function have substantial impact on achievable satisfaction rate. Furthermore, this is true for *all* considered schemes. All the other considered parameters have a much smaller impact, if any, and the relative magnitude changes between different schemes.

In our analysis we have mainly focused on schemes in which a WBSN can only pick its frequency channel, but cannot adjust its phase, i.e. the relative position in time of a BSNs periodic beacons with respect to its own time reference.[1] By comparing the considered schemes against an idealized scheme which distributes all WBSNs evenly over both frequency and time, we demonstrate that there is still a performance gap between the best frequency-adaptive scheme and the idealized scheme, which we attribute to the latter also adjusting the phases of all WBSNs meeting on the same channel. To close this gap, in future research we aim to design and evaluate a robust scheme allowing WBSNs on the same channel to negotiate their phases with the goal to minimize overlap.

To the best of our knowledge these questions have not yet received much attention as compared to the co-existence of IEEE802.15.4 with other wireless technologies operating in the 2.4 GHz band [14,15].

This paper is organized as follows: in the next Section 2 we give the necessary background on the IEEE 802.15.4 standard. Following this, in Section 3 we introduce our system model and explain the main performance measures used in this paper. The considered schemes are described in Section 4 and the sensitivity analysis is carried out in Section 5. Related work is summarized in Section 6 and we give our conclusions in Section 7.

2. Background

In this section we describe the relevant functionalities provided by the IEEE 802.15.4 standard [6].

2.1. Physical layer

The IEEE 802.15.4 standard supports different physical layers in the 2.4 GHz band. In this paper we focus on the most widely deployed one, which is the O-QPSK PHY, supporting a data rate of 250 KB/s. The 2.4 GHz band subdivided into 16 non-overlapping frequency channels. Data signals occupy 2 MHz of spectrum and the channel separation is 5 MHz. With respect to internal interference, we only consider interference from BSNs on the same channel and ignore adjacent-channel interference [16].

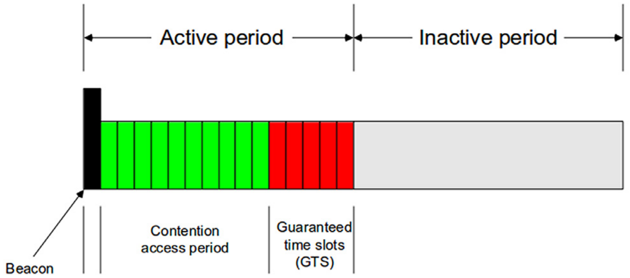

Fig. 1. Superframe structure of IEEE 802.15.4 beaconed mode.

2.2. MAC layer: beaconed mode

We assume that wireless body sensor networks use a single-hop star topology and run in the so-called beaconed mode. A star network consists of one PAN coordinator (Personal Area Network coordinator, hereafter simply called the coordinator) which starts the network and determines its major operational parameters, e.g. the frequency band, the duty cycle and others. All the other nodes (referred to as sensors or devices) first associate with the coordinator and then exchange data with it.

In the beacon-enabled mode time is sub-divided into subsequent *superframes*, which are further subdivided into an active period and an inactive period, see Fig. 1. The active period is subdivided into 16 slots. In the first slot the coordinator broadcasts a beacon frame without using carrier-sensing. Following this comes the *contention access period* (CAP), during which the devices transmit uplink packets to the coordinator or request pending downlink packets using a CSMA-type access method. Optionally, some of the 16 slots can be set apart as *guaranteed time slots*, which are allocated exclusively to individual nodes and which can be used for downlink or uplink transmissions. However, since transmissions in GTS slots are not guarded by carrier-sensing operations, GTS packets are susceptible to interference [8] and we do not consider them in this paper – for similar reasons we also disregard the ALOHA-type access method that can be alternatively used in the CAP. The sensor nodes are required to receive beacons (to maintain synchronization) and can sleep otherwise, unless they have data to transfer. The coordinator has to be switched on during the entire active period, whereas in the inactive period it can either sleep or use the time for other purposes, depending on the considered scheme (see Section 4).

The length of the superframe and the relative length of the active period within a superframe are configurable. The duration of a superframe is called “Beacon Interval” (BI) and is determined as follows:

$$BI = aBaseSuperframeDuration \times 2^{BO} \quad (1)$$

where the configurable parameter BO (“beacon order”) is an integer between 0 and 14, and $aBaseSuperframeDuration = 15.36$ ms for the 2.4 GHz O-QPSK PHY. The length of the active period is called superframe duration (SD) and is given by

$$SD = aBaseSuperframeDuration \times 2^{SO} \quad (2)$$

for $0 \leq SO \leq BO \leq 14$. The parameter SO is configurable and is called the “superframe order”.

2.3. MAC layer: network start and synchronization

The coordinator starts the BSN. In the model foreseen by the standard[2] the higher layers can instruct the MAC to scan all

---

[1] Note that the different WBSNs are not synchronized with each other and each one has its own randomly chosen phase.

[2] We are referring here to the 2011 version of the standard [6]. In 2012 the IEEE approved the IEEE 802.15.4e amendment [17] which introduces one frequency-