Computer Communications 000 (2016) 1–11

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 xxx

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].

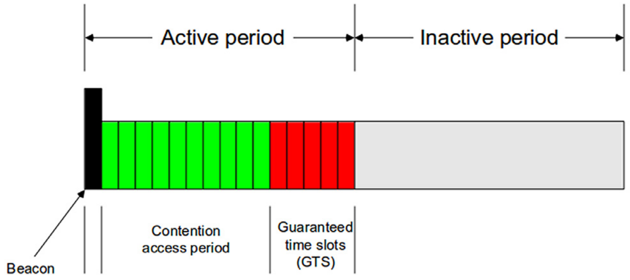

Fig. 1. Superframe structure of IEEE 802.15.4 beaconed mode.

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-

Please cite this article as: A. Moravejosharieh, A. Willig, Mutual interference in large populations of co-located IEEE 802.15.4 body sensor networks—A sensitivity analysis, Computer Communications (2016), http://dx.doi.org/10.1016/j.comcom.2016.01.002

available channels using either an energy or a passive/active MAC layer scan. The collected results can then be used by the higher layers to make decisions about the operating channel, the duty cycle and BO/SO parameters and the PAN identifier to use. Once these steps are completed, the coordinator begins to transmit beacon packets periodically.

From now on nodes are able to discover beacon packets from their coordinator. This is accomplished by scanning all frequency channels and listening on each one for a pre-determined amount of time. This discovery in itself can be a time-consuming process when the nodes know neither the frequency nor the beacon order (see [18,19]), but for the purposes of this paper we make the simplifying assumption that the sensor nodes know at least the beacon order and listen on each channel for one beacon interval before switching to the next one. After discovering a beacon packet from their pre-configured PAN coordinator, nodes attempt to associate with the coordinator by sending association request packets during the CAP.

The main purpose of beacon packet transmission is to keep the associated nodes synchronized with the coordinator, to announce the presence of downlink traffic (if any), and to specify the allocation of guaranteed timeslots (if any). After discovering the first beacon we assume that a node is required to attempt to receive all future beacon packets.[3] It will wake up shortly before expected beacon transmission in order to receive the beacon. Upon successful reception of a beacon packet, the extracted information (BO/SO values and beacon payload, if any) is delivered to the higher layers. Upon losing four consecutive beacon packets a node concludes that synchronization has been lost and informs its higher layers, which then start the searching and association process again. During the time between losing four consecutive beacons and re-discovery of its coordinator a node is said to be in the *orphan state*. In the orphan state the node cannot transmit or receive any data. Any packets generated during this time by the application on a sensor node enters a MAC buffer unless the buffer is full, then arriving packets are dropped and counted as lost.

3. System model and performance measures

In this section, the system model deployed for our simulation-based study is described and the main performance measures are defined.

3.1. System model

An individual WBSN consists of one coordinator and four sensor nodes, arranged in a star topology. Each of the four sensor nodes has a distance of 1 m to the coordinator. There is no further attenuation (e.g. coming from shadowing by the human body), and no external interference. These assumptions allows us to attribute all packet losses to packet collisions and not to path loss or hidden terminal situations.

A WBSN is operated in the beaconed mode, the BO and SO values are varied in our experiments. We consider only uplink traffic (i.e. from sensors to the coordinator), no downlink packets except acknowledgements and beacon packets are transmitted. All uplink packets are transmitted during the CAP, no GTS slots are configured.[4] The sensors first associate with the coordinator upon receiving a beacon packet with the correct MAC address. Data packets are generated periodically, the period is varied in our experiments. The payload size of data packets is fixed to 64 bytes. The coordinator responds to each successfully received data packet with an immediate acknowledgement. If a sensor does not receive an acknowledgement, it performs a bounded number of retransmissions.[5] If all retransmissions are exhausted without receiving an acknowledgement, the packet is counted as a failure, otherwise as a success. Sensor devices attempt to receive all beacons from their coordinator to stay synchronized. If the device has not received four consecutive beacon packets, it enters the orphan state. The further actions of the orphan depend on the considered scheme: in those schemes where a WBSN always operates on the same frequency the orphan node will search for beacons on the same frequency as it were before. In schemes which allow dynamic and on-going changes of the center frequency the orphan needs to assume that its WBSN might have switched to another channel and consequently has to scan all channels in round-robin fashion to rediscover its coordinator.

We assume that all WBSNs are located on the same spot. In our simulation experiments we will vary the number of WBSNs to determine one of our main performance measures (see below). We have chosen to place the WBSNs at the same spot to avoid hidden-terminal situations and to be able to explain the observed performance completely in terms of the direct internal interference experienced by WBSNs. Furthermore, this allows us to largely ignore the impact of different transmit power settings for WBSNs, and we assume that all nodes use the same transmit power. The different WBSNs are switched on at random times (except for the static-idealized scheme, see below), and there is no common time reference and no synchronization at all between different WBSNs. More specifically, for each WBSN its activation time (where both coordinator and sensors are activated simultaneously) is drawn randomly and independently from an exponential distribution with an average of one second. We also assume that the individual nodes can have clock drift. More specifically, the drift for each node is drawn from a zero-mean Gaussian with a standard deviation of 30 μs.

In Table 1 we show the values of all parameters which we have kept fixed in our study.

Table 1
Fixed parameters.

Parameter	Value
Network setup	
Layout of one WBSN	One coordinator, four sensors on a circle of 1m radius around coordinator, beaconed mode
WBSN location	All on the same spot
Number of WBSNs	All WBSNs configured identically, number varied in {50, 100, 150, 200, 250}
Channel model	Log-distance [21], no shadowing, no external interference, no hidden-terminal situations
Application Layer Parameters	
Data payload	64 byte
Coordinator start up delay	Exponential distr. 1 s
MAC Layer (CC2420) Parameters	
MAC Buffer size	16
Max. number of retransmissions	9
Physical layer (CC2420) Parameters	
Transmit power	−25 dBm
Data rate	250 kbps

---

hopping MAC mode (called TSCH) but which to the best of our knowledge has not yet found widespread deployment.

[3] This can be configured. We make this assumption to allow the higher layers to initiate operations like frequency adaptation, see below.

[4] It has been shown for example in [8] that the absence of carrier-sensing in the TDMA time slots in interference scenarios leads to substantial performance penalties in terms of packet loss rates.

[5] We have chosen not to vary the number of retransmissions, as the results reported in [20] indicate that packet success probabilities do not change significantly beyond three retransmissions, and for reliability-oriented applications we assume that one would have at least three.

Please cite this article as: A. Moravejosharieh, A. Willig, Mutual interference in large populations of co-located IEEE 802.15.4 body sensor networks—A sensitivity analysis, Computer Communications (2016), http://dx.doi.org/10.1016/j.comcom.2016.01.002

3.2. Performance measures

We use two main performance measures in this paper. They are both geared towards applications that have some notion of "acceptable" and "unacceptable" performance, e.g. in terms of packet losses for regularly transmitted sensor signals. Both are based on the notion of *satisfaction*: we regard an individual WBSN as satisfied when its average packet success rate (defined as the fraction of uplink packets generated by any sensor within the WBSN for which the originating sensor receives an acknowledgement from the coordinator, the average being taken over all four equally loaded sensors and the entire simulation time) is 95% or more. The first measure is the *satisfaction rate*, which we define as the percentage of satisfied WBSNs out of the given total number WBSNs. This total number is varied and taken from the set $\mathcal{W} = \{50, 100, 150, 200, 250\}$.

The second performance measure is the *carrying capacity*, defined as the number of WBSNs which can, under a given scheme, be located on the same spot such that the large majority of them (at least 95% of the WBSNs) are satisfied. The precise method of calculation will be explained in Section 5.4.

In addition to these two performance measures we will also show results for the average packet success rate of WBSNs, whereas above the average packet success rate of an individual WBSN is the average packet success rate of all the uplink data packets sent by the four (equally loaded) sensors of a WBSN, and the (overall) average success rate is the average of the success rates of all WBSNs. This is interesting for applications which do not have a natural threshold for acceptable packet loss performance but are able to degrade gracefully with increasing packet loss rate.

4. Considered schemes

In this section we describe the different schemes by which the PAN coordinators pick their initial frequency band and, in some schemes, change their frequency band afterwards. Here we restrict to passive schemes, i.e. schemes in which there is no active negotiation between neighbored WBSNs (and thus no exchange of control packets). We sub-divide the passive schemes into static schemes, in which a WBSN coordinator makes a decision for an operating channel only once and then never changes the channel, and dynamic schemes in which several changes of the operating channel are possible.

4.1. Static schemes

In the first baseline scheme, called the **static-random scheme**, the coordinator of a WBSN picks its frequency channel autonomously and randomly according to a uniform distribution and stays there throughout. In this scheme, no measurements are performed. An orphaned sensor node does not search through the channels but remains on the known operating channel while searching for the next beacon. Please note that the startup time and thus the phase of an individual WBSN is also chosen randomly according to the system model.

In the **static-initial-choice scheme** each WBSN coordinator scans all channels in random order at initialization time in order to estimate their load and then selects the channel with the smallest load. More precisely, a coordinator listens on each channel for one beacon period to detect as many other beacon packets as possible and then proceeds to the next channel. This procedure is repeated until all channels are covered, then the coordinator picks the channel with the fewest observed beacons (ties are broken randomly). Neither the channel nor the phase are changed after the coordinator has made its decision.

The third baseline scheme is the **static-idealized scheme** in which we assume the presence of a genie having the ability to decide both the frequency and phase of the active periods of all WBSNs. The genie uses this ability to distribute the WBSNs equally over frequency and time, i.e. such that all frequencies carry (nearly) the same number of WBSNs and on each frequency the WBSNs have an equidistant spacing in time. Clearly, as the number of WBSNs increases, there will be increased overlap of the active periods of WBSNs placed on the same channel, which might result in degradation of packet success rates. Again, in this scheme a WBSN does not switch its frequency afterwards nor does it shift its beacons / superframe phase.[6] We have introduced this scheme not only because of its fairness, but we also hypothesize that this particular allocation will give the highest average per-WBSN packet success rate and also the highest carrying capacity, and thus provides a useful yardstick for comparison.

4.2. Dynamic schemes

As a reminder, in dynamic schemes the coordinator of a WBSN is allowed to change its operating channel several times.

In the **dynamic-random-hopping scheme** the coordinator of a WBSN continuously observes the packet success rate on its current operating channel, and if it degrades below a pre-defined threshold (in this paper: the satisfaction threshold of 95%) the coordinator picks a new frequency channel randomly (with uniform distribution from all available channels except the current one) and its WBSN jumps there. The actual jump is executed in a fashion similar to [8]: the coordinator indicates the new channel in the beacon payload for four successive beacons and then jumps. To determine the success rate, the coordinator uses the sequence numbers contained in uplink packets, the size of a sequence number gap indicates the number of lost packets between two successfully received ones. In our simulations we have used a sliding window of 50 beacon periods over which we calculate the loss rate from the number of received packets and the accumulated size of sequence number gaps within this window.[7] When the coordinator decides to change the frequency, its associated sensors may become orphaned after losing 4 consecutive beacon packets, so they have to search all channels for their coordinator.

The **dynamic-targeted-hopping scheme** is similar in spirit to the lazy frequency adaptation scheme introduced in [8]. The coordinator of a WBSN uses its inactive period to continuously scan all the channels in a round-robin fashion – one channel is scanned per beacon period and the number of beacons observed during that time is counted. When the packet success rate on the current channel (obtained in the same way as for the dynamic-random-hopping scheme) drops below the 95% threshold, the coordinator decides on a new channel to operate on by choosing the one of the other 15 channels where the smallest number of beacons has been observed (ties are broken randomly).

5. Sensitivity analysis

The key goal of the sensitivity analysis carried out in this section is to explore how sensitive the satisfaction rate is to changes in a number of important system parameters, and to identify the factors having the strongest influence on the responses.

---

[6] For this scheme we assume ideal and identical clocks on all sensor nodes in this paper. In a more realistic setting clocks and thus phases would deviate over time, calling for frequent re-adjustment.

[7] This is a reasonable approach when traffic rates are reasonably high. In a more general case with larger spacing between generated packets one could use an exponentially-weighted moving-average estimator for the packet loss rate.

Please cite this article as: A. Moravejosharieh, A. Willig, Mutual interference in large populations of co-located IEEE 802.15.4 body sensor networks—A sensitivity analysis, Computer Communications (2016), http://dx.doi.org/10.1016/j.comcom.2016.01.002

Table 2
Factors, their RSM variables and their Min/Max values.

Parameter	Factor variable	Min value (-1)	Max value (1)
MAC layer (CC2420) parameters			
Beacon order	x_1	4	7
Superframe order	x_2	1	3
macMinBE	x_3	1	macMaxBE
macMaxBE	x_4	3	8
Application layer parameters			
Packet inter-arrival time	x_5	5 s	10 s

We adopt the well-known response surface methodology (RSM) [12,13,22,23]. Broadly speaking, one identifies first a desired scalar response variable (in our case: the satisfaction rate) and a number of so-called **factors**, i.e. system parameters which can be expected to have some impact on the response and which are varied. We adopt an approach where each factor takes two different values: a (sensibly defined) minimum and maximum value. With k factors a total of 2^k responses have to be obtained from simulations – this is also known as a 2^k-factorial experiment. Since an individual simulation depends on random numbers and thus generates random output, we perform 32 replications for each of the 2^k different parameter combinations and use the average of these as a response value. Each replication runs for a timespan sufficient for each sensor node to generate 10,000 packets, assuming it is associated all the time. At the end of a single replication we then calculate the packet success rate for each BSN and subsequently the number and rate of satisfied BSNs.[8] Afterwards, a regression model is fitted to the observed average responses and this regression model is then used to analyze the relative impact of the chosen factors.

We first present the factors chosen for this study (Section 5.1), next explain the RSM approach in more detail (Section 5.2) and then present our results (Section 5.4).

5.1. Factors

For our study we have chosen the following factors (see also Table 2 for the factors, their variable name in the RSM analysis and their minimal/maximal values):

- Beacon order (x_1): as discussed in Section 2.2 the beacon order (BO) parameter determines the beacon period and therefore the overall rate of beacon transmissions. We have chosen the minimum and maximum beacon orders as four and seven, respectively.
- Superframe order (x_2): the superframe order (SO) determines the time window available for sensors to send their uplink packets. We have chosen the minimum and maximum superframe orders as 1 and 3, respectively, so that each superframe order can be combined with each of the beacon orders while satisfying the constraint $SO \leq BO$.
- The macMinBE (x_3) and macMaxBE (x_4) parameters are related to the collision-avoidance CSMA MAC protocol used by IEEE 802.15.4 in the uplink: before each carrier-sensing attempt the MAC layer waits for a random backoff time. This time is a multiple of a random integer drawn uniformly from the interval $[0, 2^{BE}-1]$, where BE is the current backoff exponent. BE is initialized with *macMinBE* and increased each time the channel is sensed as busy, until the maximum value *macMaxBE* has been reached. Therefore these parameters define how aggressively a sensor accesses the channel.
- System load or packet inter-arrival time (x_5): we assume that sensors generate packets periodically with a configurable packet inter-arrival time. Please note that in general the inter-arrival time and the beacon period / beacon order are not completely independent of each other, as the beacon period must be smaller than the inter-arrival time for the latter to be meaningful. Therefore we have chosen the minimum inter-arrival time to be larger than the largest beacon period.

We argue that these factors include the most relevant MAC factors: as they determine the overall channel load generated by one WBSN and the "aggressiveness" of channel access. For all other parameters we use the default values suggested by the standard. We have also assessed the impact of some of the other parameters in preliminary studies. For example, for the number of MAC retransmissions we found that there are only minor performance differences for three or more retransmissions. It was imperative to keep the number of factors limited, as otherwise simulation times would have become prohibitive.

The RSM approach (described in Section 5.2) will allow us to obtain quantitative insight into the relative impact of these factors using only two different levels for each of them, saving many experiments as compared to a full factorial design.

5.2. RSM approach

To make the paper self-contained, we briefly summarize the RSM approach we follow in this paper, which is also known as 2^k factorial design (see also [13]). Fundamentally, in this approach the response variable Y (here: the satisfaction rate) is expressed as a function of the factors x_k, i.e. $Y = f(x_1, \ldots, x_k; \alpha)$, where α is a set of parameters for the functional form $f(\cdot)$, they are also called the regression coefficients. The parameters α are then chosen to best match the observed responses. A standard choice for $f(\cdot)$ (and the choice made in this paper) is a second-order polynomial, i.e. the response variable is expressed as

$$Y = \alpha_0 + \sum_{i=1}^{k} \alpha_i \cdot x_i + \sum_{i=1}^{k} \sum_{j<i} \alpha_{i,j} \cdot x_i \cdot x_j \qquad (3)$$

and α_0, α_i and $\alpha_{i,j}$ are the intercept, linear, and mixed coefficients (or interactions), respectively. As explained above, in the 2^k-factorial design each factor assumes either a minimum or a maximum value. To make sure that all factors enter this equation with the same order of magnitude, it is customary to not use the factors directly, but to represent the minimum value of a factor as '-1' and the maximum value as '1', i.e. we have $x_i \in \{-1, 1\}$. Furthermore, units are ignored.

For each parameter setting $(x_1, \ldots, x_k) \in \{-1, 1\}^k$ we obtain a response value $y_{x_1,\ldots,x_k}$. Observing that the regression model (Eq. 3) is linear in the parameters α_x we can represent all parameters as a vector $\alpha = (\alpha_0, \alpha_1, \ldots, \alpha_k, \alpha_{2,1}, \alpha_{3,1}, \alpha_{3,2}, \alpha_{4,1}, \ldots, \alpha_{k,k-1})$, the responses as a vector $\mathbf{y} = (y_{(-1,\ldots,-1)}, y_{(-1,\ldots,-1,1)}, \ldots, y_{(1,\ldots,1)})$ and then set up the linear equation system

$$\mathbf{y} = \mathbf{S} \cdot \alpha \qquad (4)$$

where $\mathbf{S}$ is the so-called sign matrix, in which the row corresponding to response $y_{x_1,\ldots,x_k}$ with $(x_i \in \{-1, 1\})$ is formed as $(1, x_1, \ldots, x_k, x_2 \cdot x_1, x_3 \cdot x_2, x_3 \cdot x_1, \ldots, x_k \cdot x_{k-1})$. Each matrix entry thus is either '1' or '-1'. The columns of matrix $\mathbf{S}$ are orthogonal. However, note that $\mathbf{S}$ has 2^k rows and only $1 + \frac{k(k+1)}{2}$ columns, so this linear system of equations is over-determined and we compute the least-squares solution for it.[9] The intercept α_0 is the mean

---

[8] With 32 replications we can reach a relative confidence interval half-width of 5% at a confidence level of 95% for the success rates

[9] One could make $\mathbf{S}$ quadratic, and then in fact completely orthogonal, by expanding Eq. 3 to include terms for the higher interactions of three, four, ..., k factors.

Please cite this article as: A. Moravejosharieh, A. Willig, Mutual interference in large populations of co-located IEEE 802.15.4 body sensor networks—A sensitivity analysis, Computer Communications (2016), http://dx.doi.org/10.1016/j.comcom.2016.01.002

value of all observed responses, i.e.

$$\alpha_0 = \bar{y} = \frac{1}{2^k} \sum_{x \in \{-1,1\}^k} y_x \tag{5}$$

and we use this to calculate the so-called *sum-of-squares-total* (SST)

$$SST = \sum_{x \in \{-1,1\}^k} (y_x - \bar{y})^2 \tag{6}$$

which represents the total amount of variation observed in the experiments. After elementary algebra and exploiting orthogonality of the columns of **S** one gets that

$$SST = 2^k(\alpha_1^2 + \cdots + \alpha_k^2 + \alpha_{2,1}^2 + \cdots + \alpha_{k,k-1}^2) \tag{7}$$

so that the relative impact of factor k or any interaction k, j can be expressed as:

$$\frac{2^k \alpha_k^2}{SST}, \qquad \frac{2^k \alpha_{k,j}^2}{SST} \tag{8}$$

which represents the contribution that each factor has in the total observed variation.

On the other hand, the *sum-of-squares-errors* is a measure for the total error introduced through the regression, it is given as:

$$SSE = \sum_{x \in \{-1,1\}^k} \left(y_x - \left(\alpha_0 + \sum_{i=1}^{k} \alpha_i \cdot x_i + \sum_{i=1}^{k} \sum_{j<i} \alpha_{i,j} \cdot x_i \cdot x_j \right) \right)^2 \tag{9}$$

The quality of the regression model (and here: how harmful it is to discard higher-order interactions) can be expressed as the coefficient of determination (also called the R^2 value), given by:

$$R^2 = \frac{SST - SSE}{SST} \tag{10}$$

where higher values are better.

5.3. Simulation approach

To obtain the results we have used the Castalia open-source network simulator in version 3.2, which is designed for WBSNs simulation scenarios [24], and have extended it to implement the schemes described in Section 4. The main parameters of the MAC and the packet inter-arrival times have been chosen as described in Tables 1 and 2. For the wireless channel we use the log-distance model [21]. Since the WBSNs were all placed in the same location and do not move (as described in Section 3.1) we have eliminated packet losses resulting from path loss, fading, shadowing or hidden-terminal situations, and all packet losses observed are due to direct collisions.

For each possible factor combination $c = c_{x_1,\ldots,x_5}, x_i \in \{-1, 1\}$ (i.e. each possible allocation of 1 and -1 to the five factors x_1 to x_5) and each investigated scheme we have run a number of at least 64 independent replications. Each replication lasted 10,000 or 20,000 simulated seconds, so that on average each sensor node generates 2000 packets, depending on the chosen inter-arrival time. Further replications were added when needed to achieve a relative half-width of the confidence interval not larger than 5%, at a 95% confidence level, for the success rate. From the success rates we have then determined the satisfaction rates for each replication. The satisfaction rates of all replications for one parameter allocation / scheme have then been averaged to obtain the response y_c value being used in the RSM analysis.

5.4. Results

5.4.1. Satisfaction rate

We first discuss the results for the satisfaction rate as response variable, where due to lack of space we restrict to the cases of 200 and 250 WBSNs. In Table 3 we show the results of the RSM analysis for the static-random, static-idealized, dynamic-random and dynamic-targeted schemes for the case of 200 WBSNs. The results for the static-initial-choice scheme were generally very similar to the results for the static-random scheme and are not considered furthermore. In this table we include results for the percentage contribution to variation of the linear and interaction coefficients. In Eqs. (11)–(14) we give the fitted response models (with coefficients trimmed to two decimals) for the static-idealized, static-random, dynamic-random and dynamic-targeted schemes, respectively. The same information for the case of 250 WBSNs is shown in Table 4 and in Eqs. (15)–(18). Furthermore, in Fig. 2 we plot for these four schemes the percentage impact of the linear terms for each of the considered variables and for all considered numbers of WBSNs, and in Table 5 we report the intercept values (α_0) for all considered schemes and different numbers of WBSNs. Recall from Table 2 that factor x_1 is the beacon order, factor x_2 the superframe order, x_3 is

Table 3
Main RSM results for the satisfaction rates of all considered schemes for 200 WBSNs.

	Static-idealized	Static-random	Dynamic-random	Dynamic-targeted
SST	59342.9	36979.0	59159.2	59651.2
SSE	1324.6	2038.0	1368.8	1014.4
R^2	0.98	0.94	0.98	0.98
α_0	58.28	26.16	44.25	54.28
% contrib. x_1	1.41	4.66	0.52	1.43
% contrib. x_2	1.8	0.25	7.0e−2	1.15
% contrib. x_3	87.71	57.81	89.44	90.81
% contrib. x_4	1.04	10.41	2.96	1.31
% contrib. x_5	0.5	2.95	1.71	0.5
% contrib. x_1x_2	0.38	4.0e−2	0.0	0.4
% contrib. x_1x_3	1.34	4.85	0.78	0.77
% contrib. x_1x_4	4.0e−2	7.0e−2	0.46	2.0e−2
% contrib. x_1x_5	0.0	1.71	8.0e−2	2.0e−2
% contrib. x_2x_3	1.32	0.33	0.11	0.76
% contrib. x_2x_4	0.36	7.0e−2	8.0e−2	0.3
% contrib. x_2x_5	0.42	3.0e−2	7.0e−2	0.34
% contrib. x_3x_4	0.49	9.96	1.38	0.19
% contrib. x_3x_5	0.29	1.81	0.29	4.0e−2
% contrib. x_4x_5	0.0	1.27	0.45	0.0

Table 4
Main RSM results for the satisfaction rates of all considered schemes for 250 WBSNs.

	Static-idealized	Static-random	Dynamic-random	Dynamic-targeted
SST	54563.8	29760.2	49048.3	53443.7
SSE	1769.3	1823.8	1949.2	1825.4
R^2	0.97	0.94	0.96	0.97
α_0	38.75	20.15	33.01	37.09
% contrib. x_1	2.13	10.8	2.1	2.0
% contrib. x_2	0.0	0.61	3.0e−2	4.0e−2
% contrib. x_3	75.74	43.23	68.82	75.38
% contrib. x_4	7.06	10.26	9.72	6.86
% contrib. x_5	1.56	3.02	1.98	1.83
% contrib. x_1x_2	0.0	8.0e−2	3.0e−2	6.0e−2
% contrib. x_1x_3	2.11	10.79	2.09	2.5
% contrib. x_1x_4	1.77	0.86	0.7	1.18
% contrib. x_1x_5	5.0e−2	2.2	0.34	9.0e−2
% contrib. x_2x_3	3.0e−2	0.65	0.0	0.0
% contrib. x_2x_4	1.0e−2	0.27	0.0	0.0
% contrib. x_2x_5	1.0e−2	9.0e−2	1.0e−2	0.1
% contrib. x_3x_4	5.89	10.09	8.94	5.59
% contrib. x_3x_5	0.45	2.19	0.87	0.71
% contrib. x_4x_5	0.83	0.74	1.43	1.32

Please cite this article as: A. Moravejosharieh, A. Willig, Mutual interference in large populations of co-located IEEE 802.15.4 body sensor networks—A sensitivity analysis, Computer Communications (2016), http://dx.doi.org/10.1016/j.comcom.2016.01.002

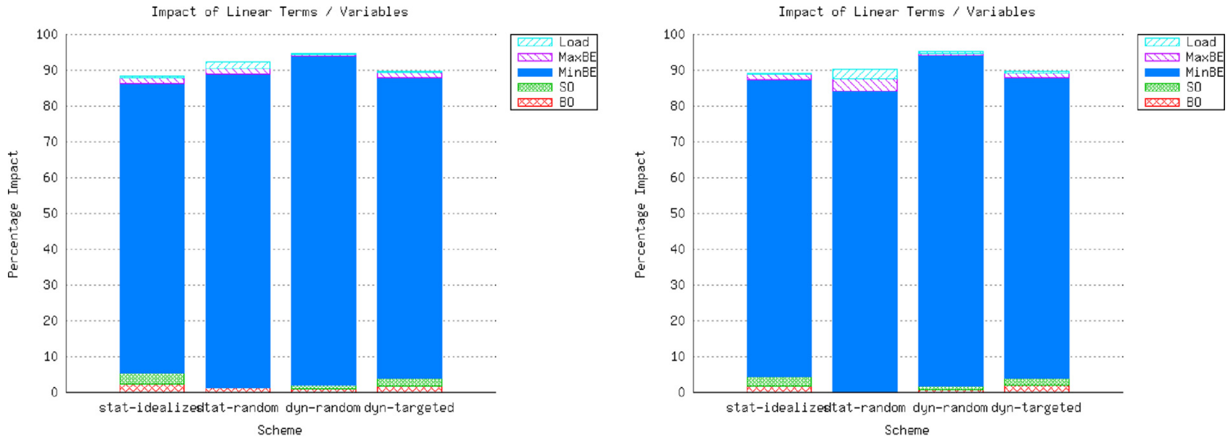

(a) Impact of factors on satisfaction rate, WBSN-density = 50

(b) Impact of factors on satisfaction rate, WBSN-density = 100

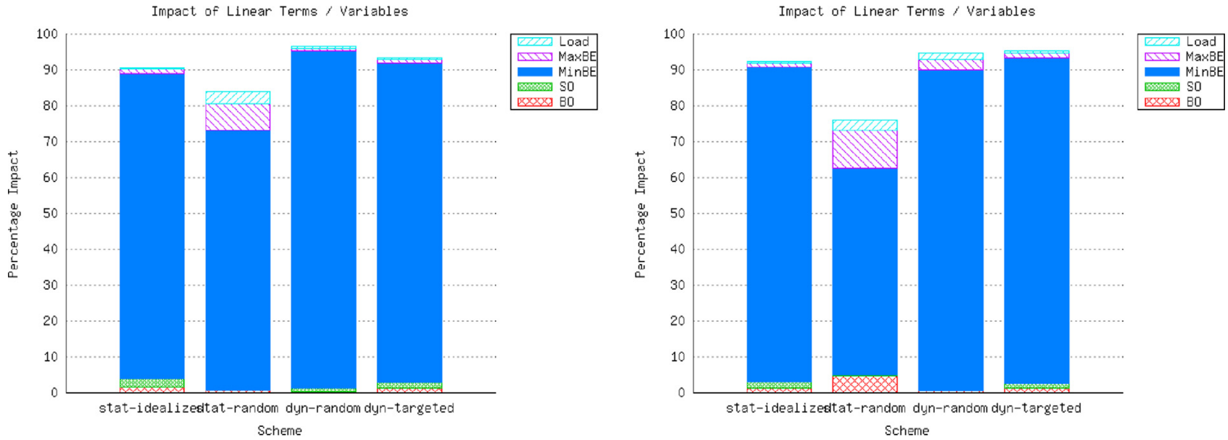

(c) Impact of factors on satisfaction rate, WBSN-density = 150

(d) Impact of factors on satisfaction rate, WBSN-density = 200

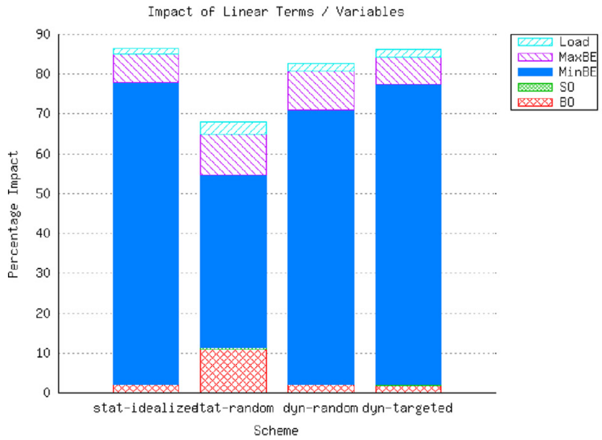

(e) Impact of factors on satisfaction rate, WBSN-density = 250

Fig. 2. Percentage impacts of different factors on satisfaction rate in the presence of varying internal interference.

macMinBE, x_4 is *macMaxBE* and x_5 is the packet inter-arrival time.

$$\begin{aligned} Y_{\text{static-idealized}} &= 58.28 \\ &-5.11 \cdot x_1 - 5.77 \cdot x_2 + 40.33 \cdot x_3 + 4.39 \cdot x_4 + 3.03 \cdot x_5 \\ &+2.64 \cdot x_1x_2 + 4.99 \cdot x_1x_3 - 0.82 \cdot x_1x_4 - 0.16 \cdot x_1x_5 + 4.96 \cdot x_2x_3 \\ &-2.57 \cdot x_2x_4 - 2.80 \cdot x_2x_5 - 3.00 \cdot x_3x_4 - 2.31 \cdot x_3x_5 + 0.29 \cdot x_4x_5 \end{aligned} \tag{11}$$

$$\begin{aligned} Y_{\text{static-random}} &= 26.16 \\ &+ 7.33 \cdot x_1 + 1.70 \cdot x_2 + 25.85 \cdot x_3 + 10.97 \cdot x_4 + 5.83 \cdot x_5 \\ &+ 0.67 \cdot x_1x_2 + 7.49 \cdot x_1x_3 + 0.87 \cdot x_1x_4 + 4.45 \cdot x_1x_5 \\ &+ 1.96 \cdot x_2x_3 + 0.93 \cdot x_2x_4 - 0.63 \cdot x_2x_5 + 10.73 \cdot x_3x_4 \\ &+ 4.57 \cdot x_3x_5 - 3.83 \cdot x_4x_5 \end{aligned} \tag{12}$$

Please cite this article as: A. Moravejosharieh, A. Willig, Mutual interference in large populations of co-located IEEE 802.15.4 body sensor networks—A sensitivity analysis, Computer Communications (2016), http://dx.doi.org/10.1016/j.comcom.2016.01.002

Table 5
Intercept values for the satisfaction rate for all considered schemes and different numbers of WBSNs

# WBSNs	Static-idealized	Static-random	Dynamic-random	Dynamic-targeted
50	62.11	46.75	56.77	59.44
100	61.01	41.26	54.57	58.71
150	59.76	33.14	52.08	56.34
200	58.28	26.16	44.25	54.28
250	38.75	20.15	33.01	37.09

$$Y_{\text{dynamic-random}} = 44.25 +3.09 \cdot x_1 - 1.17 \cdot x_2 + 40.66 \cdot x_3 + 7.40 \cdot x_4 + 5.63 \cdot x_5 +0.10 \cdot x_1x_2 + 3.80 \cdot x_1x_3 - 2.91 \cdot x_1x_4 - 1.24 \cdot x_1x_5 + 1.44 \cdot x_2x_3 -1.21 \cdot x_2x_4 - 1.13 \cdot x_2x_5 + 5.05 \cdot x_3x_4 + 2.33 \cdot x_3x_5 - 2.89 \cdot x_4x_5 \tag{13}$$

$$Y_{\text{dynamic-targeted}} = 54.28 -5.16 \cdot x_1 - 4.62 \cdot x_2 + 41.14 \cdot x_3 + 4.94 \cdot x_4 + 3.05 \cdot x_5 +2.73 \cdot x_1x_2 + 3.80 \cdot x_1x_3 - 0.63 \cdot x_1x_4 + 0.57 \cdot x_1x_5 + 3.77 \cdot x_2x_3 -2.37 \cdot x_2x_4 - 2.53 \cdot x_2x_5 - 1.89 \cdot x_3x_4 - 0.89 \cdot x_3x_5 - 0.27 \cdot x_4x_5 \tag{14}$$

$$Y_{\text{static-idealized}} = 38.75 +6.03 \cdot x_1 - 0.09 \cdot x_2 + 35.94 \cdot x_3 + 10.97 \cdot x_4 + 5.16 \cdot x_5 +0.23 \cdot x_1x_2 + 6.00 \cdot x_1x_3 - 5.50 \cdot x_1x_4 + 0.90 \cdot x_1x_5 + 0.75 \cdot x_2x_3 -0.36 \cdot x_2x_4 - 0.29 \cdot x_2x_5 + 10.02 \cdot x_3x_4 + 2.76 \cdot x_3x_5 - 3.76 \cdot x_4x_5 \tag{15}$$

$$Y_{\text{static-random}} = 20.15 +10.02 \cdot x_1 + 2.38 \cdot x_2 + 20.05 \cdot x_3 + 9.77 \cdot x_4 + 5.30 \cdot x_5 +0.88 \cdot x_1x_2 + 10.02 \cdot x_1x_3 + 2.83 \cdot x_1x_4 + 4.52 \cdot x_1x_5 + 2.46 \cdot x_2x_3 +1.60 \cdot x_2x_4 - 0.93 \cdot x_2x_5 + 9.69 \cdot x_3x_4 + 4.52 \cdot x_3x_5 - 2.63 \cdot x_4x_5 \tag{16}$$

$$Y_{\text{dynamic-random}} = 33.01 +5.67 \cdot x_1 - 0.72 \cdot x_2 + 32.48 \cdot x_3 + 12.21 \cdot x_4 + 5.51 \cdot x_5 +0.70 \cdot x_1x_2 + 5.65 \cdot x_1x_3 - 3.27 \cdot x_1x_4 + 2.29 \cdot x_1x_5 - 0.22 \cdot x_2x_3 -0.22 \cdot x_2x_4 - 0.46 \cdot x_2x_5 + 11.70 \cdot x_3x_4 + 3.65 \cdot x_3x_5 - 4.68 \cdot x_4x_5 \tag{17}$$

$$Y_{\text{dynamic-targeted}} = 37.09 +5.77 \cdot x_1 - 0.79 \cdot x_2 + 35.48 \cdot x_3 + 10.71 \cdot x_4 + 5.53 \cdot x_5 +0.96 \cdot x_1x_2 + 6.47 \cdot x_1x_3 - 4.44 \cdot x_1x_4 + 1.21 \cdot x_1x_5 + 0.08 \cdot x_2x_3 +0.10 \cdot x_2x_4 - 1.32 \cdot x_2x_5 + 9.67 \cdot x_3x_4 + 3.45 \cdot x_3x_5 - 4.70 \cdot x_4x_5 \tag{18}$$

The presented results highlight a number of interesting points:

- The R^2 values for all four schemes and for both presented numbers of WBSNs (Tables 3, 4 for 200 and 250 WBSNs, the values for other numbers are similar) indicate that the regression ansatz given in Eq. 3 explains already almost all the observed variation. This means that the incomplete model (in which we ignored all interactions between three or more factors) is a good approximation.
- The static-idealized and dynamic-targeted schemes have achieved the highest average satisfaction rate α_0 for both 200 and 250 WBSNs. For these two schemes it can be seen from Fig. 2 that the linear terms (i.e. the factors) alone already explain more than 80% of the observed variation (for 200 and fewer WBSNs: more than 90%). Furthermore, both these schemes are similar in the sense that the factor x_3 (*macMinBE*) is by far the most influential one and for the case of 250 WBSNs the factor x_4 (*macMaxBE*) has also noticeable impact, whereas the impact of the other factors is negligible. For both the dynamic-targeted-hopping and static-idealized schemes with 250 WBSNs the x_3x_4 interaction explains most of the remaining variation not covered by the linear terms. These findings, together with the observation that for all schemes the factors x_3 (*macMinBE*) and x_4 (*macMaxBE*) are the most dominant ones (compare Fig. 2) suggests that *these two variables together have the most impact on satisfaction rate performance (with macMinBE being the much more important one), and the adaptation or at least careful configuration of these parameters offers substantial potential for improvement.*
- When judging from the intercept values (see Table 5), the static-idealized scheme is the overall best scheme. The dynamic-targeted-hopping scheme is the second best, but comes relatively close to the static-idealized scheme. However, for all but three of the 32 parameter combinations the static-idealized scheme outperforms the dynamic-targeted-hopping scheme. In the static-idealized scheme the BSNs are evenly distributed in both frequency and time / phase, whereas the dynamic-targeted scheme only achieves an equal distribution of BSNs across frequencies but sticks to the initial random distribution of phases. This suggests then that *the performance difference between the static-idealized and the dynamic-targeted scheme can be explained by the lack of phase adjustment in the latter scheme, and at the same time the size of the performance gap is an indication of what can be gained by adjusting the phases as well.*
- Again judging from the intercept values, the dynamic-random scheme shows modest performance differences to the dynamic-targeted-hopping scheme, but is still substantially better than the static-random scheme. *This suggests that indeed frequency hopping can provide substantial benefits in terms of satisfaction rate.*
- The static-random scheme also stands out in how the relative impact of the factors changes when the number of WBSNs is increased (compare Fig. 2): on the one hand the impact *macMinBE* factor becomes smaller much more rapidly than for other schemes, and at the same time the impact of the factor x_1 (*BO*) grows the largest for increasing number of WBSNs.

Clearly, our results indicate that the *macMinBE* parameter (and to a lesser extent *macMaxBE*) plays a decisive role in the achievable performance. From the regression equations (Eqs. 11–18) the term for x_3 (*macMinBE*) enters with a positive sign, so to achieve better satisfaction rate we need to choose larger values of *macMinBE*. The *macMinBE* parameter determines the initial average waiting time after which a sensor node performs a carrier-sense operation for a new packet, so it is a measure for how aggressively a node tries to send data. Longer initial waiting times lead to fewer collisions so that more channel resources are left for useful transmissions even in high node densities.

5.4.2. Carrying capacity

In Table 6 we show the results for the carrying capacity of the different schemes (see Section 3.2). To obtain the carrying capacity, we simulate a given scheme with WBSN numbers taken from the set $\mathcal{W}$ so that for each of these numbers $w \in \mathcal{W}$ a number

Please cite this article as: A. Moravejosharieh, A. Willig, Mutual interference in large populations of co-located IEEE 802.15.4 body sensor networks—A sensitivity analysis, Computer Communications (2016), http://dx.doi.org/10.1016/j.comcom.2016.01.002

Table 6
Carrying capacity.

Schemes	Carrying capacity
Static-random	19
Static-initial-choice	20
Dynamic-random-hopping	93
Dynamic-targeted-hopping	137
Static-idealised	155

Table 7
Intercept values for the packet success rate for all considered schemes and different numbers of WBSNs.

# WBSNs	Static-random	Static-idealized	Dynamic-random	Dynamic-targeted
50	85.95	89.64	86.84	88.46
100	84.37	88.21	85.28	86.81
150	81.85	86.4	82.86	84.75
200	79.14	84.36	80.64	82.49
250	72.31	78.92	74.1	76.5

of 64 replications is carried out. For each replication we calculate the number of satisfied WBSNs, and we compute the average of these numbers over all replications, giving us the average percentage of satisfied WBSNs for a given number $w \in \mathcal{W}$ of WBSNs. After collecting these averages for all WBSN numbers from $\mathcal{W}$ we calculate a regression curve (a second-order polynomial) allowing to interpolate the average number of satisfied WBSNs between the points given by $\mathcal{W}$, and the carrying capacity is determined as the point / number of WBSNs where this regression curve crosses the 95% line. Please note that we have resorted to this interpolation approach since otherwise simulation times would have been prohibitively long. We show the carrying capacity for the following parameters: $BO = 64$, $SO = 4$, packet inter-arrival time of one second, $macMinBE = 3$, and $macMaxBE = 5$ (the latter two are the default values suggested by the standard). It is interesting to note that here the difference between the dynamic-targeted-hopping and the static-idealized scheme is more pronounced. Again, we essentially attribute the difference between these two schemes to the inability of the dynamic-targeted-hopping scheme to adjust the phases of the WBSNs sharing the same channel.

5.4.3. *Packet success rate*

The main focus of this paper is on performance measures geared towards applications having some notion of “acceptable” and “unacceptable” packet loss performance and using a particular threshold to distinguish between these (here we use a packet success rate of 95% to mark a WBSN as satisfied). For applications which do not have a natural threshold it is interesting to get some insight into the packet success rate performance of the different schemes. We have again carried out a RSM analysis of all four schemes for the packet success rate (see Section 3.2). In Table 7 we show for all considered numbers of WBSNs and all considered schemes the intercept values α_0 for the packet success rate. The following points are interesting:

- While not shown here, again the chosen second-order model has a very high R^2 value (≥ 0.98) for all schemes and all numbers of WBSNs, so it explains almost all of the observed variation.
- The intercept values α_0 of the different schemes are closer to each other than we have found for the satisfaction rate, and for the same scheme their range is relatively smaller. This suggests that introducing sharp thresholds makes differences between the schemes more pronounced.

Table 8
Intercept values for the packet success rate standard deviation for all considered schemes and different numbers of WBSNs.

# WBSNs	Static-random	Static-idealized	Dynamic-random	Dynamic-targeted
50	4.02	3.67	4.06	3.92
100	4.35	3.78	4.26	4.04
150	4.63	3.51	4.4	3.99
200	4.99	3.96	4.86	4.5
250	5.9	4.29	5.5	4.79

Another important aspect of the packet success rate is how different it can be for different WBSNs in the same scenario, i.e. how fair the packet success rate allocation to WBSNs is. To look into this we have carried out the RSM analysis for the average standard deviations of the packet success rates. The intercept values α_0 for this are shown in Table 8. It can be seen that generally the average standard deviation does not exceed 6%, which is relatively small compared to the average packet success rate percentages reported in Table 7, and suggests that on average the differences in packet success rates among WBSNs are minor. This can be attributed to the fact that all WBSNs are configured in the same way.

6. Related work

The performance of IEEE 802.15.4 and its dependence on individual parameters has been assessed in a wide range of studies. In [25–27] the performance of slotted CSMA/CA and the impact of various beacon order values on throughput, average delay and successful transmissions has been studied, and in [28] the authors evaluated the performance of the unbeaconed 802.15.4 MAC as the number of sensor networks increases. The impact of the backoff exponent parameters (in particular *macMinBE*) and payload size have been studied in [29,30] for IEEE 802.15.4 and Zigbee. Anastasi et al have investigated the effects of changing parameters like the packet arrival pattern (Poisson or periodic), the CSMA/CA parameters, beacon interval, packet size and different offered loads [31]. The impact of system parameters such as the packet arrival rate, number of stations, and packet size on the medium access probability or the queue length distribution is studied in [32] using Markov chain and queueing models. A detailed analytical evaluation of the CSMA/CA performance in IEEE 802.15.4 is carried out in [33], where the authors presented a Markov model that predicts the behavior of the slotted CSMA/CA mechanism in IEEE 802.15.4. The results were compared to simulation results to verify the accuracy of the developed model. Their analysis was inspired by [34] and [35] for the usage of a per-user Markov model. In [44,45] the effects of superframe overlaps and beacon collisions on co-located IEEE 802.15.4 networks in beaconed mode (used for example as BSNs or in industrial automation) have been investigated in both static and mobile scenarios. The results confirm that superframe overlaps (and to a lesser extent direct beacon collisions) lead to significant goodput losses, as nodes from different WBSNs compete for the channel. Our work differs from these papers not only by considering a wider range of parameters, but also by assessing their relative impact in scenarios with over-subscribed channel resources. We have also chosen performance measures which in our view better reflect the considered scenarios.

In our own previous work we have studied the problem of finding the potential white spaces in the spectrum without using central coordination or any kind of infrastructure [11,36,37]. We have studied the channel utilization as the network density increases for a range of schemes, including the frequency-adaptation schemes considered in this paper. The frequency adaptation feature enables WBSNs to switch to a channel with supposedly lower interference,

Please cite this article as: A. Moravejosharieh, A. Willig, Mutual interference in large populations of co-located IEEE 802.15.4 body sensor networks—A sensitivity analysis, Computer Communications (2016), http://dx.doi.org/10.1016/j.comcom.2016.01.002

but comes at the cost of increased energy consumption at the coordinator.

There has also been a range of other studies investigating adaptation schemes in the presence of interference. The authors of [38] propose to adjust the values of the BO and SO parameters under heavy traffic loads. Another algorithm for adjusting the beacon order is proposed in [39] with the aim to avoid collisions, to improve channel utilization and reduce energy consumption when the traffic load of the network is high. The adaptive adjustment of the beacon interval is also proposed in [40] where the main aim is to prolong the lifetime of a sensor network. In some of the referenced papers, the IEEE MAC protocol has been modified noticeably to avoid internal interference. In some cases the considered number of WBSN was relatively small, in other cases the improved performance gains came with the higher costs of embedding a new module and higher complexity. While our paper lays a foundation for the design of mitigation strategies for internal interference, the authors of [46] address the problem of classifying (external) interference from within an IEEE 802.15.4 sensor network, which is a fundamental building block towards the design of a suite of specifically tailored mitigation strategies for external interference. The particular approach taken in the paper analyzes individual received packets for “fingerprints”, e.g. signal-strength or LQI variations over received packets that are typical for different types of external interferers (Bluetooth, WiFi and others).

The IEEE 802.15.6 standard [41] (approved in 2012) has introduced four strategies to mitigate the interference caused by neighboring BANs: beacon-shifting, channel-hopping, active superframe interleaving and B2-aided time-shifting (only offered in the non-beacon enabled mode). For example, the beacon-shifting approach allows WBSNs to postpone their beacons to avoid active period overlapping. However, Kim et al. have proposed a flexible beacon scheduling scheme where coordinators perform carrier sensing before beacon transmission. They have compared their approach to beacon-shifting and found significant improvements over the latter [42]. As the number of occupants of the channel increases, the probability of flexible beacon scheduling failure increases as well, since sensing a busy channel leads coordinators to back off instead of sending the beacon packets. Bradai et al have conducted a simulation study to compare the performance of IEEE 802.15.4 and IEEE 802.15.6 [43] and have observed that IEEE 802.15.4 achieves higher throughput than IEEE 802.15.6 for increasing number of sensor nodes.

7. Conclusions

In this paper we have considered a scenario where a large number of WBSNs meets at one point and competes for channel resources (time and frequency). In particular, we have assessed the relative impact of a range of static and adaptive schemes for frequency allocation on two new major performance metrics, the satisfaction rate and the carrying capacity. We have found that for most schemes by far the most influential parameters are related to the MAC backoff process. By comparing the considered frequency-adaptive schemes against an idealized scheme allocating WBSNs in both frequency and time, we have furthermore found that on the one hand frequency adaptation can provide substantial gains, but on the other hand further gains can potentially be reached by adding the capability to properly separate WBSNs operating on the same channel in time (“phase adjustment”).

There is substantial potential for future work. As a next step, we plan to design phase adjustment schemes which operate without any centralized entity and are integrated with frequency adaptation. Secondly, we intend to assess the considered schemes in more general scenarios with dynamic WBSN populations, the presence of hidden-terminal situations (i.e. where not all WBSNs are located on the same spot) and heterogeneous loads. We are also currently working on a measurement campaign to gain experimental insight into the performance of the schemes discussed in this paper. Another interesting direction is to extend this work towards a combination of internal and external interference, e.g. coming from WiFi or Bluetooth.

References

[1] J. Penders, B. Gyselinckx, R. Vullers, M. De Nil, V. Nimmala, J. van de Molengraft, F. Yazicioglu, T. Torfs, V. Leonov, P. Merken, C.V. Hoof, Human++: from technology to emerging health monitoring concepts, in: Proceedings of the 2008 5th International Summer School and Symposium on Medical Devices and Biosensors, 2008, pp. 94–98.
[2] J. Penders, J. van de Molengraft, L. Brown, B. Grundlehner, B. Gyselinckx, C.V. Hoof, Potential and challenges of body area networks for personal health, Annual International Conference of the IEEE Engineering in Medicine and Biology Society, 2009. EMBC 2009., IEEE, 2009, pp. 6569–6572.
[3] S. Ullah, H. Higgins, B. Braem, B. Latre, C. Blondia, I. Moerman, S. Saleem, Z. Rahman, K.S. Kwak, A comprehensive survey of wireless body area networks, J. Med. Syst. 36 (3) (2012) 1065–1094.
[4] H. Cao, V. Leung, C. Chow, H. Chan, Enabling technologies for wireless body area networks: a survey and outlook, Commun. Mag., IEEE 47 (12) (2009) 84–93.
[5] P.S. Hall, Y. Hao, Y.I. Nechayev, A. Alomainy, C.C. Constantinou, C. Parini, M.R. Kamarudin, T.Z. Salim, D.T. Hee, R. Dubrovka, Antennas and propagation for on-body communication systems, Antennas Propag. Mag., IEEE 49 (3) (2007) 41–58.
[6] 2011, LAN/MAN Standards Committee of the IEEE Computer Society, IEEE Standard for Local and metropolitan area networks – Part 15.4: Low Rate Wireless Personal Area Networks (LR-WPANs), June 2011, revision of.
[7] J. Zheng, M. Lee, Will ieee 802.15.4 make ubiquitous networking a reality?: a discussion on a potential low power, low bit rate standard, Commun. Mag., IEEE 42 (6) (2004) 140–146.
[8] E.T. Yazdi, A. Willig, K. Pawlikowski, Frequency adaptation for interference mitigation in IEEE 802.15.4-based mobile body sensor networks, Elsevier Comput. Commun. 53 (2014) 102–119.
[9] R. Musaloiu-E, A. Terzis, Minimising the effect of wifi interference in 802.15.4 wireless sensor networks, Int. J. Sens. Netw. 3 (1) (2008) 43–54.
[10] J.-H. Hauer, V. Handziski, A. Wolisz, Experimental study of the impact of wlan interference on ieee 802.15.4 body area networks, in: Proceedings of the 6th European Conference on Wireless Sensor Networks (EWSN), Cork, Ireland, 2009.
[11] A. Moravejosharieh, E.T. Yazdi, A. Willig, Study of resource utilization in ieee 802.15.4 wireless body sensor network, part ii: greedy channel utilization, Proceedings of the 19th IEEE International Conference on Networks (ICON), 2013, IEEE, 2013, pp. 1–6.
[12] R. Myers, D. Montgomery, Response Surface Methodology, John Wiley and Sons,Inc, 2002.
[13] R. Jain, The Art of Computer Systems Performance Analysis – Techniques for Experimental Design, Measurement, Simulation, and Modeling, Wiley Professional Computing, John Wiley and Sons, New York, Chichester, 1991.
[14] W. Yuan, X. Wang, J.-P. Linnartz, A coexistence model of ieee 802.15. 4 and ieee 802.11 b/g, Proceedings of the 14th IEEE Symposium on Communications and Vehicular Technology in the Benelux, 2007, IEEE, 2007, pp. 1–5.
[15] H. Huo, Y. Xu, C.C. Bilen, H. Zhang, Coexistence issues of 2.4 ghz sensor networks with other rf devices at home, Third International Conference on Sensor Technologies and Applications, 2009. SENSORCOMM’09., IEEE, 2009, pp. 200–205.
[16] E. Toscano, L.L. Bello, Cross-channel interference in ieee 802.15.4 networks, Proceedings of the 7th IEEE International Workshop on Factory Communication Systems (WFCS), Dresden, Germany, 2008.
[17] IEEE Computer Society, IEEE standard for local and metropolitan area networks – part 15.4: Low-rate wireless personal area networks (LR-WPANs) – amendment 1: MAC sublayer, 2012, IEEE Std 802.15.4e-2012.
[18] A. Willig, N. Karowski, J.-H. Hauer, Passive discovery of IEEE 802.15.4-based body sensor networks, Ad Hoc Netw. 8 (7) (2010) 742–754.
[19] N. Karowski, A.C. Viana, A. Wolisz, Optimized asynchronous multi-channel discovery of ieee 802.15.4-based wireless personal area networks, IEEE Trans. Mobile Comput. 12 (10) (2013) 1972–1985.
[20] D. Rohm, M. Goyal, H. Hosseini, A. Divjak, Y. Bashir, Configuring beaconless ieee 802.15. 4 networks under different traffic loads, Proceedings of the International Conference on Advanced Information Networking and Applications, 2009. AINA’09., IEEE, 2009, pp. 921–928.
[21] T.S. Rappaport, Wireless Communications – Principles and Practice, Prentice Hall, Upper Saddle River, NJ, USA, 2002.
[22] L. Wang, T.J. Kazmierski, J. Wenninger, Response-surface-based design space exploration and optimisation of wireless sensor nodes with tunable energy harvesters, Proceedings of the Conference on Design, Automation and Test in Europe, EDA Consortium, 2012, pp. 733–738.
[23] R. Unal, R.A. Lepsch, M.L. McMillin, Response surface model building and multidisciplinary optimization using d-optimal designs, in: Proceedings of the 7th AIAA/USAF/NASA/ISSMO Symposium on multidisciplinary Analysis and optimisation, 1998, pp. 405–411.

Please cite this article as: A. Moravejosharieh, A. Willig, Mutual interference in large populations of co-located IEEE 802.15.4 body sensor networks—A sensitivity analysis, Computer Communications (2016), http://dx.doi.org/10.1016/j.comcom.2016.01.002

[24] Castalia, A simulator for WSNs, [Online]. Available: http://castalia.npc.nicta.com.au/.
[25] A. Koubaa, M. Alves, E. Tovar, A comprehensive simulation study of slotted csma/ca for ieee 802.15. 4 wireless sensor networks, IEEE WFCS 6 (2006) 63–70.
[26] S. Pollin, M. Ergen, S. Ergen, B. Bougard, L.D. Perre, I. Moerman, A. Bahai, P. Varaiya, F. Catthoor, Performance analysis of slotted carrier sense ieee 802.15.4 medium access layer, IEEE Trans. Wirel. Commun. 7 (9) (2008) 3359–3371.
[27] P. Pelegris, K. Banitsas, Investigating the efficiency of ieee 802.15.4 for medical monitoring applications, in: Proceedings of the Annual International Conference of the IEEE Engineering in Medicine and Biology Society, EMBC, 2011, 2011, pp. 8215–8218.
[28] N. Golmie, D. Cypher, O. Rebala, Performance analysis of low rate wireless technologies for medical applications, Comput. Commun. 28 (10) (2005) 1266–1275. Performance issues of Wireless LANs, {PANs} and ad hoc networks
[29] Z. Tao, S. Panwar, D. Gu, J. Zhang, Performance analysis and a proposed improvement for the ieee 802.15.4 contention access period, Wireless Communications and Networking Conference, 2006. WCNC 2006. IEEE, 4, IEEE, 2006, pp. 1811–1818.
[30] J.S. Choi, M. Zhou, Performance analysis of zigbee-based body sensor networks, IEEE International Conference on Systems Man and Cybernetics (SMC), 2010, IEEE, 2010, pp. 2427–2433.
[31] G. Anastasi, M. Conti, M.D. Francesco, A comprehensive analysis of the mac unreliability problem in ieee 802.15. 4 wireless sensor networks, IEEE Trans. Ind. Inf. 7 (1) (2011) 52–65.
[32] J. Mii, S. Shafi, V.B. Mii, The impact of mac parameters on the performance of 802.15.4 pan, Ad Hoc Netw. 3 (5) (2005) 509–528. Data Communication and Topology Control in Ad Hoc Networks.
[33] S. Pollin, M. Ergen, S. Ergen, B. Bougard, L.D. Perre, I. Moerman, A. Bahai, P. Varaiya, F. Catthoor, Performance analysis of slotted carrier sense ieee 802.15.4 medium access layer, IEEE Trans. Wirel. Commun. 7 (9) (2008) 3359–3371.
[34] G. Bianchi, Performance analysis of the ieee 802.11 distributed coordination function, IEEE J. Select. Areas Commun. 18 (3) (2000) 535–547.
[35] I. C. S. L. M. S. Committee, Wireless lan medium access control (mac) and physical layer (phy) specifications, 1997.
[36] A. Moravejosharieh, E. Yazdi, A. Willig, K. Pawlikowski, Adaptive channel utilisation in IEEE 802.15.4 wireless body sensor networks: continuous hopping approach, in: Proceedings of the Telecommunication Networks and Applications Conference (ATNAC), 2014 Australasian, 2014, pp. 93–98.
[37] A. Moravejosharieh, A. Willig, K. Pawlikowski, Frequency-adaptive approach in IEEE 802.15.4 wireless body sensor networks:continuous-assessment or periodic-assessment? Australas. J. Inf., Commun. Technol. Appl. 1 (1) (2015).
[38] R. Sari, Analysis of the effect of beacon order and superframe order value to the performance of multihop wireless networks on ieee 802.15. 4 protocol, International Conference on Advanced Computer Science and Information Systems (ICACSIS), 2012, IEEE, 2012, pp. 89–94.
[39] M. Neugebauer, J. Plönnigs, K. Kabitzsch, A new beacon order adaptation algorithm for ieee 802.15. 4 networks, Proceedings of the Second European Workshop on Wireless Sensor Networks, 2005., IEEE, 2005, pp. 302–311.
[40] L. Qiang, T. Jun, Minimum-energy-cost algorithm based on superframe adaptation control, in: IEEE International Conference on Communications (ICC), 2011, 2011, pp. 1–5.
[41] Ieee standard for local and metropolitan area networks - part 15.6: Wireless body area networks, 2012, IEEE Std 802.15.6-2012, pp. 1–271.
[42] S. Kim, S. Kim, J.-W. Kim, D.-S. Eom, Flexible beacon scheduling scheme for interference mitigation in body sensor networks, Proceedings of the 9th Annual IEEE Communications Society Conference on Sensor, Mesh and Ad Hoc Communications and Networks (SECON), 2012, IEEE, 2012, pp. 157–164.
[43] N. Bradai, L. Fourati, L. Kamoun, Performance analysis of medium access control protocol for wireless body area networks, in: Proceedings of the 27th International Conference on Advanced Information Networking and Applications Workshops (WAINA), 2013, 2013, pp. 916–921.
[44] N. Nordin, F. Dressler, Effects and implications of beacon collisions in co-located IEEE 802.15.4 networks, in: Proceedings of the IEEE Vehicular Technology Conference (VTC Fall), Quebec, Canada, 2012, pp. 1–5.
[45] , On the impact of beacon collisions in co-located IEEE 802.15.4-based networks, in: Proceedings of the IEEE INFOCOM Student Workshop, Orlando, Florida, 2012.
[46] F. Hermans, O. Rensfelt, T. Voigt, E. Ngai, L.-A. Norden, P. Gunningberg, SoNIC: classifying interference in 802.15.4 sensor networks, in: Proceedings of the 12th IEEE/ACM Conference on Information Processing in Sensor Networks (IPSN), Philadelphia, Pennsylvania, 2013.

Please cite this article as: A. Moravejosharieh, A. Willig, Mutual interference in large populations of co-located IEEE 802.15.4 body sensor networks—A sensitivity analysis, Computer Communications (2016), http://dx.doi.org/10.1016/j.comcom.2016.01.002