Passive discovery of IEEE 802.15.4-based Body Sensor Networks☆

Andreas Willig*, Niels Karowski, Jan-Hinrich Hauer,☆☆

Abstract

In this paper we study passive discovery of IEEE 802.15.4 networks operating in the beacon-enabled mode. The task of discovery occurs in different scenarios. One example is a simple device that wishes to associate with a specific, pre-specified PAN coordinator (targeted discovery). Another example are opportunistic relaying applications, where arbitrary foreign coordinators should be discovered (untargeted discovery). We consider a simple class of listening strategies and provide different analytical models which allow to find optimal strategies for different listening scenarios. For the case of targeted discovery without constraints on the listening costs we give a dynamic programming formulation, for targeted discovery with bounded costs we present and validate a simple model and derive the desired performance measures. For untargeted discovery we present simulation results in a mobile scenario.

Key words: IEEE 802.15.4, passive discovery, stochastic modeling

☆We gratefully acknowledge the partial support of this research activity by the European project FP6-2005-IST-5-033506 ANGEL

☆☆All authors are with the Telecommunication Networks Group (TKN) at Technische Universität Berlin, Germany

*Corresponding authors address: Andreas Willig, Fachbgebiet Telekommunikationsnetze, Sekr. FT-5, Einsteinufer 25, 10587 Berlin, Germany, email: awillig@tkn.tu-berlin.de

Email address: {awillig,karowski,hauer}@tkn.tu-berlin.de (Andreas Willig, Niels Karowski, Jan-Hinrich Hauer,☆☆)

Corresponding authors address:	Andreas Willig Telecommunication Networks Group (TKN) Technische Universität Berlin Sekr. FT-5 Einsteinufer 25 10587 Berlin Germany
email:	awillig@ieee.org
fon:	+49 - 30 - 314 23836
fax:	+49 - 30 - 314 23818

1. Introduction

Body sensor networks (BSN) are expected to play a major role in future health- and wellness-related services and systems [1], [2]. They give people the freedom to move around while their vital functions are monitored and diagnosed. Body sensor networks have some similarities with "normal" fixed wireless sensor networks [3], but there are also important differences, for example the comparably small number of nodes and the specific mobility pattern (group mobility at pedestrian speeds). The small geographical size of body sensor networks makes IEEE 802.15.4 [4] personal area networks (PAN) a particularly attractive networking technology [5], [6].

One fundamental task that shows up in various forms is to *discover* an IEEE 802.15.4 network. As an example, consider a simple sensor node (a reduced function device in IEEE 802.15-4-parlance) being statically configured to join one specific BSN, identified for example by the unique address of the network coordinator. The node should join exactly the specified BSN and no other, so that one persons medical data are not mistaken for another persons data. A second example, in which nodes of one BSN must detect a *foreign* BSN (or foreign PAN, FPAN), are schemes for opportunistic message relaying [7], [8], [9], [10], [11], [12], [13]. In these schemes the mobility of BSNs is exploited, together with data replication between BSNs, to disseminate data. From now on, we refer to both the specific coordinator in the first example and the arbitrary coordinator in the second example as an FPAN.

The IEEE 802.15.4 standard offers a physical layer operating in the 2.4 GHz ISM band. This band is subdivided into sixteen orthogonal channels. An IEEE 802.15.4 network selects one of them and operates there all the time (static frequency allocation). To discover another network means to determine its major communication parameters, including its frequency, duty cycle and beacon period. The applicable discovery strategies depend on the type of the network to be discovered. In the so-called unbeaconed mode of IEEE 802.15.4 it is possible to use *active discovery*: the discovering device (henceforth called the *listener*) sends a request packet and the PAN coordinator answers with a response packet. If no response is received, the listener switches to the next channel until all channels are exhausted. This procedure is straightforward and quick, but requires separate signaling packets. In the so-called beaconed-mode of IEEE 802.15.4 only passive discovery is available. Here, the listener is not allowed to send extra request packets but has to listen for the occurence of periodic beacon packets from the coordinator to be discovered. The beaconed mode is suitable whenever a certain quality-of-service is required in a network, since it allows to partially run the network in TDMA mode. Once the listener receives a beacon, the communication parameters of the FPAN are known. Of course, there is a desire to organize the passive listening process such that FPANs are reliably detected and the listening times are small. The major control knob for this is the *listening schedule*. Such a schedule determines when to listen on which channel and for how long.

The goal of this paper is to identify good listening schedules for different scenarios. In one scenario a specific, a-priori known FPAN has to be detected. In another scenario we have opportunistic relaying application in mind and consider detection of an arbitrary FPAN. We consider a relatively simple class of listening schedules, the so-called *sweep strategies*. We have made this choice for two reasons. The first reason was simply to keep the problem tractable. The second reason addresses scenarios in which the initiation of a listening process and the actual execution of the process are separated. More specifically, a coordinator should be able to instruct one or more of its members to perform listening on behalf of him. However, the coordinator must specify the listening strategy to the listener, and the *description length* of the strategy becomes an issue. The sweep strategies can be chosen to have comparably small description lengths. Therefore, they require only small communication overhead, or, in cases where the strategies are pre-deployed on the nodes, a small amount of memory. For the class of sweep strategies we present models allowing to compute optimal strategies under certain constraints.

The paper is structured as follows: in the next Section 2 we provide the necessary background on the IEEE 802.15.4 standard and its discovery facilities. In Section 3 we describe the system model and the considered class of sweep strategies. Following this, in Section 4 we consider a specific discovery scenario called unbounded targeted discovery, in which the listener searches for a specific PAN without any limits on listening time. For this scenario a dynamic-programming model is adopted. In Section 5 we address targeted discovery with a limited time-budget and provide an analytical model for the discovery probability and listening times. In Section 6 we discuss two further variations of the discovery problem: the usage of multiple listeners and the search for arbitrary FPANs in a mobile scenario. Related work is reviewed in Section 7 and our conclusions are given in Section 8.

2. Background

The IEEE 802.15.4 low-rate wireless personal area network (LR-WPAN) standard [4] was finalized in October 2003, a revised version has been published at the end of 2006 (we use the terms PAN and BSN interchangeably). It covers the physical layer and the MAC layer in the OSI model.

2.1. Channelization and node types

An IEEE 802.15.4 node can work in one of 27 frequency channels, placed in three different frequency bands: there is one channel in the range from 868 to 868.6 MHz, ten channels in the range from 902 to 928 MHz, and 16 channels in the 2.4 GHz ISM bands. Typically, an IEEE 802.15.4 network selects one of these channels at its discretion and stays on it. In this paper we concentrate on the 2.4 GHz PHY.

The standard defines different types of nodes with different responsibilities: *full-function devices* (FFD) and *reduced-function devices* (RFD). RFDs implement only a subset of the full protocol functionality. A full-function device can

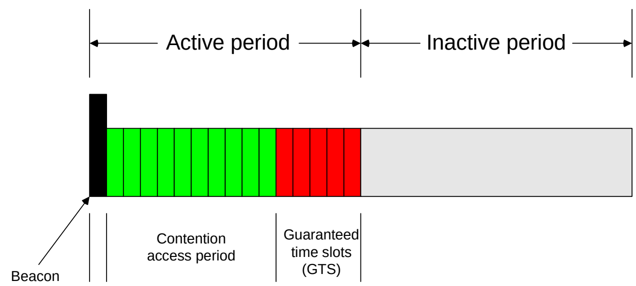

Figure 1: Superframe structure of IEEE 802.15.4

operate in three different roles: as a *PAN coordinator* (or network coordinator), as a *coordinator*, or as a *device*. In contrast, a RFD can only act in the role of a device. There is only a single PAN coordinator in a network, but there could be several coordinators. The PAN coordinator starts the network and selects the major operational parameters, including the PAN identifier, the frequency channel and the duty cycle (see below). Simple devices have to associate to exactly one coordinator and can only exchange packets with it, but not with other devices. A coordinator and its associated devices form a star network.

2.2. The beacon-enabled mode

The protocol offers two different modes: the *non-beacon-enabled mode* and the *beacon-enabled mode*. In the beacon-enabled mode the time is subdivided into consecutive *superframes*, the structure of a superframe is shown in Figure 1. A superframe is subdivided into an active period and an inactive period. At the beginning of the active period the coordinator broadcasts a *beacon packet*, but without performing a carrier-sense operation. The beacons therefore occur in a strictly periodic fashion. The length of the superframe and the relative length of the active period within a superframe (the *duty cycle*) are configurable. More specifically, the superframe length and therefore the beacon period is given by [4, Sec. 7.5.1.1]:

$$aBaseSuperframeDuration \cdot 2^{BO}$$

where $aBaseSuperframeDuration = 15.36$ ms (for the 2.4 GHz PHY) and $BO \in \{0, 1, \ldots, 14\}$ is the configurable *beacon order*. The duration of the active period is given by

$$aBaseSuperframeDuration \cdot 2^{SO}$$

with $0 \leq SO \leq BO \leq 14$ and where SO is the configurable *superframe order*. Therefore, the allowed beacon periods are restricted to powers of two times a base constant.

During the inactive period all nodes, including the coordinator, can sleep. The active period is subdivided into 16 slots, the beacon packet is always transmitted at the beginning of the first slot. The beacon packet contains, among other things, the communication parameters (BO, SO) selected for this BSN. At the end of the active period a maximum of seven *guaranteed time slots* (GTS) can be allocated to nodes in an exclusive manner. In the remaining slots (called *contention access period*, CAP) the associated nodes can send uplink packets to the coordinator or they can request pending data from the coordinator. During this time they compete for the medium using a slotted CSMA-scheme.

2.3. IEEE 802.15.4 discovery support

Different kinds of discovery support are available. In non-beacon-enabled networks an active discovery scheme is used. Here, the searching device broadcasts a request frame and coordinators respond with beacon frames. The beacon-enabled mode, however, is limited to passive discovery [4, Sec. 7.5.2.1.2], in which the searching device must listen on the medium for the (periodic) beacon packets of surrounding coordinators.

To discover the presence or absence of BSNs, the MAC management service provides a primitive `MLME-SCAN` that initiates a channel scan over a given list of channels. Two different passive scanning techniques are available: the passive scan and the energy detection scan. The *energy detection scan* allows to obtain the maximum detected energy for each requested channel, without giving any indication about the identity or type of the radiating entity. In the *passive scan*, a device only listens for beacons on the requested channels without transmitting beacon requests. The passive scan reports back detected beacons. Alternatively, a device can enable the *promiscuous mode*, in which the radio is switched to receive mode. The usual address filtering mechanism is disabled and all subsequently received frames (including packets from different networks) are delivered to the higher layers.

The standard does not support delegation of discovery tasks from, say, a coordinator to one or more helper nodes. Without such a delegation mechanism a coordinator wishing to transmit a message into a foreign BSN would have to listen itself, and during this listening time it cannot provide services to its own network. This delegation of listening (and data transfer) tasks is a key feature of a protocol developed by us and briefly mentioned in Section 6.1.

3. System model and considered listening strategies

3.1. System model

To avoid notational clutter, we describe the assumptions for the case of a single listener wishing to find beacons of a single FPAN coordinator.

The listener works in slotted time. A time slot has the duration of $aBaseSuperframeDuration = 15.36$ ms. It corresponds to the smallest beacon period with beacon order $BO = 0$. The time the listener spends on one particular channel is always an integer multiple of the slot time, fractional times are not considered.[1] The listener starts its search at time $t = 0$. A foreign BSN might or might not be present (i.e. within reception range of the listener). If it is present, then it operates on frequency channel $F \in \mathcal{F} = \{1, 2, \ldots, F_{\max}\}$, where $F_{\max}$ is the maximum allowed channel. Without further restrictions we have $F_{\max} = 16$ for the 2.4 GHz PHY and we will assume this for the rest of the paper.[2] We assume that F is drawn randomly from $\mathcal{F}$ according to a uniform distribution. The foreign BSN operates with a beacon order $B \in \mathcal{B} = \{0, 1, \ldots, B_{\max}\}$ where $B_{\max}$ is the maximal allowed beacon order. Without further restrictions we have $B_{\max} = 14$ and we assume this for the rest of this paper. The actual beacon order B is drawn randomly from $\mathcal{B}$ with probability mass function $p_B(\cdot)$. The beacon period of the foreign BSN is thus 2^B slots. Since the foreign BSN might have been started at some random time in the past, we assume that its beacons have a certain *phase shift* Φ with respect to time $t = 0$. This phase shift depends on the foreign BSNs beacon order B and is assumed to be uniformly distributed over the interval $\left\{0, 1, \ldots, 2^B - 1\right\}$. The random variables F and B are independent of each other, the phase shift Φ is independent of F and conditionally independent of B. The realizations of F, B and Φ are not known to the listener. However, the listener knows the probability distributions for F and B. The parameters are summarized in Table 1.

We make the worst-case assumption that the foreign BSN only transmits beacons, i.e. there are no other packets which would allow the listener to detect its presence.[3]

3.2. Major discovery types and performance figures

In the remaining paper we distinguish among the following different types of discovery: targeted and untargeted discovery.

The goal in *targeted discovery* is to discover one specific and pre-configured coordinator. This task occurs for example before a simple sensor device can associate to a pre-configured PAN coordinator. The coordinator is assumed to be present. We can distinguish two-cases:

- In *unbounded targeted discovery* there are no bounds on the allowable time for discovery, i.e. the listener searches until the coordinator is found. Assuming that the coordinator is discovered eventually, the primary goal is then to minimize the (average) listing time before the coordinator is found.

[1] This restriction is made for two reasons. Firstly, by this discretization we can achieve relatively easy computational procedures. Secondly, when spending less than one slot time on a channel, the overhead constituted by the channel switching times (which for the popular ChipCon CC2420 transceiver are around 1ms [14]) starts to become significant.

[2] Our models can be applied in a very similar fashion in the 902 MHz band as well. In the 868 MHz band with its single channel the task of listening becomes trivial.

[3] Of course, when the listener picks up any other packet from the FPANs coordinator or one of its members, it has identified the right channel and needs only wait for the next beacon.

Symbol	Meaning
$F_{\max}$	maximum number of frequencies used
$\mathcal{F} = \{1, 2, \ldots, F_{\max}\}$	the set of all frequencies
F	random variable denoting actual frequency of foreign BSN, drawn uniformly from $\mathcal{F}$
$B_{\max}$	maximum beacon order used
$\mathcal{B} = \{0, 1, \ldots, B_{\max}\}$	the set of possible beacon orders
B	random variable denoting actual beacon order of foreign BSN, drawn from $\mathcal{B}$ according to the pmf $p_B(\cdot)$
$W_{\max}$	maximum random waiting time
Φ	phase shift of foreign BSN w.r.t. $t = 0$, drawn uniformly from $[0, 2^B]$

Table 1: Summary of all symbols used in this paper

- In *bounded targeted discovery* the listener listens at most for a given time budget, afterwards it gives up. In this scenario there are two performance measures that need to be balanced: the probability of successful discovery, and the (average) time until discovery.

The goal in *untargeted discovery* is to find one or more arbitrary FPAN coordinators. This is relevant for example in opportunistic message relaying applications. We can again distinguish two cases. In *unbounded untargeted discovery* the listening time is not bounded, and the most interesting performance measure is the average time until the first FPAN is detected. In *bounded untargeted discovery* the listening time is bounded.

3.3. The sweep strategies

The considered class of listening strategies is called *sweep strategies*. They are quite straightforward and have been chosen for their simplicity and the fact that they have a short description length (measured in the required number of parameters to describe a strategy).

The basic unit of the sweep strategies is a *sweep* of a given *sweep order* s: In a sweep of order s the listener listens subsequently on all channels in $\mathcal{F}$, starting from channel one. On each channel the listener listens for a contiguous time of s slots, then the next channel is visited. A sweep strategy $\mathcal{S} = (s_1, s_2, \ldots, s_k)$ consists of k subsequent sweeps, with the i-th sweep having order s_i (i.e. spending s_i slots on each channel). Furthermore, between two successive sweeps an

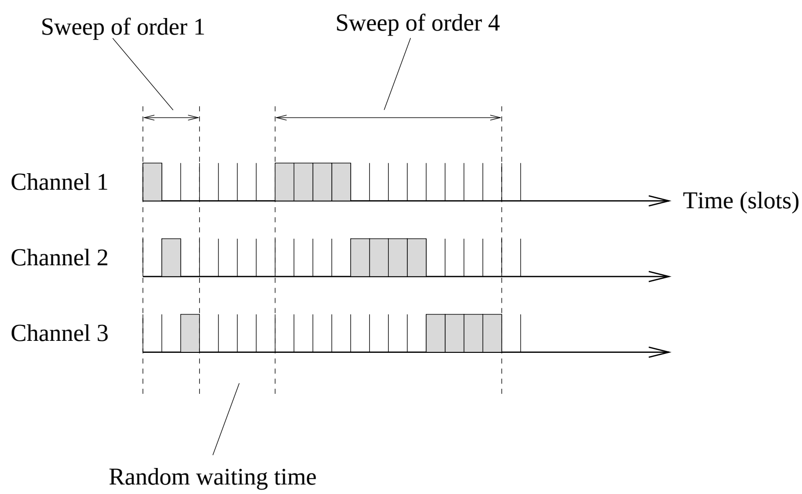

Figure 2: Example sweep strategy: $F_{\max} = 3$, one sweep of order one followed by one sweep of order four, $\mathcal{S} = (1, 4)$.

additional random waiting time is inserted, which is drawn from a uniform distribution over $[0, W_{\max}]$. In Section 5 we shed some light on the role of this random waiting time. When targeted discovery is used, the listener stops after having found the target coordinator or, in case of bounded targeted discovery, when the allowed total listening time has been exhausted.

As an example, in Figure 2 a setup with three available channels ($F_{\max} = 3$) is shown, where two sweeps are performed, one of order $s_1 = 1$ and one of order $s_2 = 4$. Although shown otherwise in the figure, it is not necessary that the random waiting time assumes an integral number of slots.

4. Unbounded targeted discovery: a Dynamic-Programming formulation

4.1. Basics of dynamic programming

The framework of (stochastic) dynamic programming (DP) [15], [16] offers a generic approach towards finding optimal strategies for controlling a dynamical system with minimum costs or maximum reward. We consider the discrete-time case only. A system at time $k \in \mathbb{N}_0$ is characterized by its current *state* x_k, which belongs to a given state space $\mathcal{X}$. In state x_k at time k the controller can choose a *control action* a_k from an action space $\mathcal{A}$, taking knowledge of the current state x_k and possibly all previous states and previous actions into account. Having chosen action a_k, the system evolves according to

$$x_{k+1} = f(x_k, a_k, w_k)$$

where w_k is a random "disturbance" and $f(\cdot)$ is a function determining the next state out of the current state, the chosen action and the disturbance. Furthermore, when at time k and in state x_k the action a_k is chosen, a cost $g_k(x_k, a_k)$ is incurred. The goal of dynamic programming is to find a sequence of policies $\mu_0(\cdot), \mu_1(\cdot), \ldots$, so that $\mu_k(\cdot)$ maps at time k the state x_k to an action a_k and the total (expected) incurred cost $E\left[\sum_{k=0}^{\infty} g_k(x_k, \mu_k(x_k))\right]$ is minimized, subject to the state evolution equation $x_{k+1} = f(x_k, \mu_k(x_k), w_k)$.

4.2. A search model with an infinite number of stages

We consider a DP formulation with an infinite number of stages (i.e. there is no bound for the time parameter k), and we assume that the FPAN is indeed present. We furthermore assume that the channel is noisy and packets may get lost. More specifically, beacon packets are lost independently of each other, each one with packet loss probability $r \in [0, 1]$.

The DP model adopted here is a variation of a model for optimal search given in [15, Sec. 3.5], the derivation of the main result follows the lines given there.

The state of the system is given by a probability vector $\mathbf{P} = (P_0, P_1, \ldots, P_{14})$, where P_b denotes the probability that the FPAN has beacon order $B = b$ (in fact, $\mathbf{P}$ is time-dependent, but we suppress this to avoid notational clutter). The available actions are sweeps of order $s \in \mathcal{A} = \{1, 2, \ldots, s_{\max}\}$, where $s_{\max}$ is a system parameter. When the FPAN has beacon order $B = b$ and the listener listens for a sweep of order s, then the probability that the FPAN is discovered is given by:

$$\alpha(s, b) = 1 - r^{\left\lfloor \frac{s}{2^b} \right\rfloor} \cdot \left(1 - (1 - r)\left(\frac{s}{2^b} - \left\lfloor \frac{s}{2^b} \right\rfloor\right)\right) \tag{1}$$

for all admissible values of s and b. This can be seen as follows. The expression $k^* = \left\lfloor \frac{s}{2^b} \right\rfloor$ gives the number of beacons that are guaranteed to occur during s time slots. In fact, during time s there can either occur k^* beacons (with probability $\omega_1 = 1 - \frac{s - k^* 2^b}{2^b}$) or $k^* + 1$ beacons (with probability $\omega_2 = \frac{s - k^* 2^b}{2^b}$). In the first case, to fail beacon reception the listener must fail all k^* trials, which happens with probability r^{k^*}, in the second case complete failure occurs with probability r^{k^*+1}. Using the law of total probability and simplifying $\omega_1 \cdot r^{k^*} + \omega_2 \cdot r^{k^*+1}$ gives the result. For the special case $r = 0$ and $s < 2^b$ (i.e. $k^* = 0$) we make the assumption that $0^0 = 1$. Please observe that for fixed b the function $s \mapsto \alpha(s, b)$ is continuous, piecewise linear (with pieces of length 2^b), and concave.

With this, we can express for a given probability distribution $\mathbf{P}$ for the FPANs beacon order, and when listening with a sweep of order s, the overall probability to discover the FPAN as:

$$\alpha(s, \mathbf{P}) = \sum_{b=0}^{B_{\max}} P_b \cdot \alpha(s, b) \tag{2}$$

Now suppose that the listener did not find the FPAN after a sweep of order s. This observation actually modifies the initial probabilities P_b that the FPAN has $B = b$. The posteriori probability vector $T_s(\mathbf{P})$ for this case is given by:

$$[\mathbf{T}_s(\mathbf{P})]_b = \frac{P_b \cdot (1 - \alpha(s, b))}{1 - \alpha(s, \mathbf{P})}$$

as can be seen from an application of Bayes' rule. After a failed sweep of order s the state is updated with these posterior probabilities. Now, be s and t two actions. It is straightforward to verify the relationship

$$(1 - \alpha(s, T_t(\mathbf{P}))) \cdot (1 - \alpha(t, \mathbf{P})) = (1 - \alpha(t, T_s(\mathbf{P}))) \cdot (1 - \alpha(s, \mathbf{P})) \tag{3}$$

which intuitively says that the probability of not finding the FPAN after looking first for s and then for t is the same as if one first looks for t and then for s. With this relationship it is then easy to see that:

$$T_t(T_s(\mathbf{P})) = T_s(T_t(\mathbf{P})) \tag{4}$$

holds.

We assume that choosing action s incurs a cost $g(s)$, where $g(\cdot)$ is assumed to be a non-negative function of s only, not of the current state $\mathbf{P}$.[4] A search policy $\pi = (s_1, s_2, \ldots)$ with $s_i \in \mathcal{A}$ first uses a sweep of order s_1. If the FPAN is not found, it next uses a sweep of order s_2 and so on. Let (s, t, π) denote a policy which first searches with order s, then with order t and then continues according to policy π. The quantity $V_\pi(\mathbf{P})$ denotes the average search cost for policy π for successfully finding the foreign PAN when the start state is $\mathbf{P}$.

Lemma 1. *We have:*

$$V_{(s,t,\pi)}(\mathbf{P}) \leq V_{(t,s,\pi)}(\mathbf{P}) \Longleftrightarrow \frac{\alpha(s, \mathbf{P})}{g(s)} \geq \frac{\alpha(t, \mathbf{P})}{g(t)}$$

Proof. The average cost of the strategy (s, t, π) is given by:

$$V_{(s,t,\pi)}(\mathbf{P}) = g(s) + (1 - \alpha(s, \mathbf{P})) \cdot [g(t) + (1 - \alpha(t, T_s(\mathbf{P}))) \cdot V]$$

where the abbreviation $V := V_\pi(T_t(T_s(\mathbf{P}))) = V_\pi(T_s(T_t(\mathbf{P})))$ is used (compare Equation 4). This equation expresses that we first have to spend costs $g(s)$ for the search with sweep order s. If this sweep fails (with probability $1 - \alpha(s, \mathbf{P})$) we spend costs $g(t)$ for a sweep of order t. If this fails again (with probability $(1 - \alpha(t, T_s(\mathbf{P}))))$ the costs V for policy π are incurred. Similarly:

$$V_{(t,s,\pi)}(\mathbf{P}) = g(t) + (1 - \alpha(t, \mathbf{P})) \cdot [g(s) + (1 - \alpha(s, T_t(\mathbf{P}))) \cdot V]$$

Simplifying the relation

$$V_{(s,t,\pi)}(\mathbf{P}) \leq V_{(t,s,\pi)}(\mathbf{P})$$

and considering Equation 3 yields the result. □

[4]This implies that a sweep is fully executed, even if the foreign PAN is found. This simplification is made to keep the model tractable.

PER	Optimal strategy	Avg. Listening time
$r = 0$	$(1, 2, 4, 8, 16, 32, 64, 128, 256, 512, 1024, 2048, 4096, 8192, 16384, 16384, \dots)$	34423.7
$r = 0.1$	$(2, 4, 8, 16, 32, 32, 64, 64, 128, 128, 256, 256, 512, 512, 512, 1024, 1024, 1024, 2048, 2048, 2048, 4096, 4096, 4096, 8192, 8192, 8192, 8192, 8192, 16384, \dots)$	67858.6
$r = 0.2$	$(2, 4, 8, 8, 16, 16, 32, 32, 64, 64, 128, 128, 128, 256, 256, 256, 512, 512, 512, 1024, 1024, 1024, 1024, 2048, 2048, 2048, 2048, 4096, 4096, 4096, 4096, 4096, 8192, 8192, 8192, 8192, 8192, 16384, \dots)$	84165.7

Table 2: Optimal strategies for the infinite-stages DP algorithm for initial state probability vector $\mathbf{P} = (1/15, 1/15, \dots, 1/15)$, action space $\mathcal{A} = \{1, 2, \dots, 2^{14}\}$, cost function $g(s) = F_{\max} \cdot s + 16$ and different values for the packet error probability r

With this lemma in place, the optimal policy can be stated as follows:

Theorem 1. *In state* $\mathbf{P}$ *an optimal policy starts with a sweep of order* s *that maximizes the value* $\frac{\alpha(s,\mathbf{P})}{g(s)}$ *among all* $s \in \mathcal{A}$.

The proof is omitted, since the proof given in [15, Sec. 3.5, Proposition 2] directly carries over to the present situation. It is essentially based on an exchange argument, repeatedly invoking Lemma 1 and showing that an optimal policy for state $\mathbf{P}$ starts with the action s maximizing $\frac{\alpha(s,\mathbf{P})}{g(s)}$. With this result, it is possible to compute the optimal policy for a given starting vector $\mathbf{P}_0$. First, one determines the optimal action s_0 for $\mathbf{P}_0$, and subsequently computes the posterior vector $\mathbf{P}_1 = T_{s_0}(\mathbf{P}_0)$. In the second round, one computes the optimal action s_1 for $\mathbf{P}_1$, computes the posterior $\mathbf{P}_2 = T_{s_1}(\mathbf{P}_1)$ and so on.

We present two numerical examples. In both examples the initial state vector is chosen as $\mathbf{P}_0 = (1/15, \dots, 1/15)$ to reflect maximum uncertainty about the FPANs beacon order. All computations have been done with arbitrary-precision rational numbers, so the given results are exact. In the first study the action space has been set as $\mathcal{A} = \{1, 2, 3, \dots, 2^{14}\}$, the cost function has been fixed as $g(s) = s \cdot F_{\max} + 16$, corresponding to a maximum random waiting time of $W_{\max} = 32$ slots after a sweep. For this setup, in Table 2 the optimal policies for different packet error rates r and the achieved expected listening times until success (the latter obtained from simulation, see Section 5.2 for a brief description of the simulation approach) are shown. The following remarkable observations can be made:

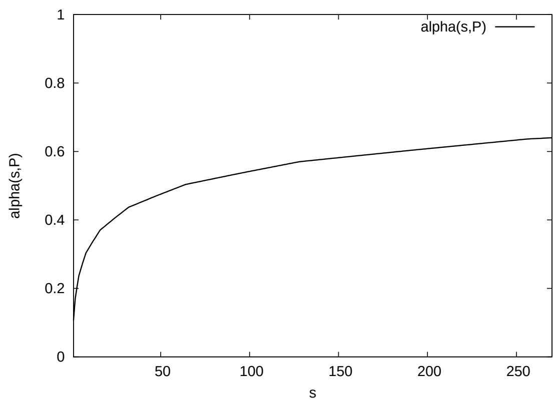

Figure 3: Graph of function $\alpha(s, \mathbf{P})$ for $\mathbf{P} = (1/15, 1/15, \ldots, 1/15)$ and PER $r = 0.2$

- Even for error-free channels ($r = 0$) it takes a very long time on average to discover the coordinator – the 34,423.7 time slots given in Table 2 correspond to approximately nine minutes. The presence of packet errors can significantly increase the average listening times, with 10% errors the average listening time is almost doubled.
- The optimal strategies start with small sweep orders and steeply increase the sweep orders after subsequent failures. With channel errors the same sweep order is used a number of times, but after that again there is is a steep increase.
- The sweep orders are all powers of two. An intuitive explanation is that the beacon periods of the FPAN are also powers of two and sweep orders not being equal to a power of two therefore include wasted time. For a more detailed explanation see below.

The preference for powers of two can be understood by looking at $\alpha(s, \mathbf{P})$ for one particular $\mathbf{P}$. In Figure 3 the curve $\alpha(s, \mathbf{P})$ is shown for one fixed vector $\mathbf{P}$. It appears to be (but is not exactly) a piecewise linear curve. However, the piecewise linear approximation $\alpha^*(s, \mathbf{P})$ (for which the values of $\alpha(s, \mathbf{P})$ for $s = 2^k$ have been chosen as anchor points) is fairly accurate.[5] And, considering

[5] Because of the law of total probability it suffices to consider the difference between $\alpha(s, b)$ and its linear approximation $\alpha^*(s, b)$. Suppose that s is chosen such that $2^a < s < 2^{a+1}$, then $\alpha^*(s, b) = \alpha(2^a, b) + \frac{s-2^a}{2^a}\left(\alpha(2^{a+1}, b) - \alpha(2^a, b)\right)$. For $b \geq a$ it can be verified that $\alpha(s, b) - \alpha^*(s, b) = 0$ holds, for $b = a - 1$ an upper bound of $|\alpha(s, b) - \alpha^*(s, b)| \leq r^2 \cdot (1 - r)^2$ can be established, and for $b = a - k$ for some $k \geq 2$ an upper bound of $2 \cdot r^{2^k} \cdot (1 - r)$ can be verified. With these bounds the linearization approximates $\alpha(\cdot)$ tightly for the considered values of the packet error probability r and probability vector $\mathbf{P}$.

Cost function	Optimal strategy
$g(s) = s \cdot F_{\max}$	$(1, 1, 1, 1, 1, \ldots)$
$g(s) = s \cdot F_{\max} + 1$	$(1, 2, 4, 8, 16, 32, 64, 128, 256, 512, \ldots)$
$g(s) = s \cdot F_{\max} + 16$	$(1, 2, 4, 8, 16, 32, 64, 128, 256, 512, \ldots)$
$g(s) = s \cdot F_{\max} + 17$	$(2, 4, 8, 16, 32, 64, 128, 256, 512, \ldots)$
$g(s) = s \cdot F_{\max} + 38$	$(2, 4, 8, 16, 32, 64, 128, 256, 512, \ldots)$
$g(s) = s \cdot F_{\max} + 39$	$(2, 8, 16, 32, 64, 128, 256, 512, \ldots)$

Table 3: Optimal strategies for the infinite-stages DP algorithm for initial state probability vector $\mathbf{P} = (1/15, 1/15, \ldots, 1/15)$, action space $\mathcal{A} = \{1, 2, \ldots, 512\}$, packet error probability $r = 0$ and different cost functions

the chosen cost function $g(s) = s \cdot F_{\max} + 16$, it is clear that for the linear approximation within each segment $[2^i, 2^{i+1}]$ the maximum of $\frac{\alpha^*(s,\mathbf{P})}{g(s)}$ can only be found in $s = 2^i$.

In the second example we look at the influence of the cost function. More specifically, we have used a cost function of the form $g(s) = s \cdot F_{\max} + C$ where $C \geq 0$ is some constant. For the results shown in Table 3 the action space has been chosen as $\mathcal{A} = \{1, 2, \ldots, 512\}$ and the packet error probability has been fixed as $r = 0$. It can be seen that the optimal policy indeed depends on the "fixed costs" C. With $C = 0$ the optimal policy uses, for the given initial probability vector $\mathbf{P}$, sweeps of length 1 all the time, for $C > 0$ there always exists a point beyond which action 1 is no longer optimal and increasingly larger sweeps (all powers of two) are used.

To summarize, the optimal search strategies for unbounded targeted discovery have, for the particular initial state probability vector $\mathbf{P} = (1/15, \ldots, 1/15)$ chosen here, two interesting properties. First, they use only powers of two as sweep orders, and secondly, they gradually increase the sweep orders with a speed depending on the channel error probability r and the considered cost function. However, the latter behaviour depends on the initial state probability vector $\mathbf{P}$: for a uniform vector it indeed makes sense to first try the smaller values (where with small investments already a significant part of the set of all outcomes can be covered) and to gradually increase the sweep orders. For $\mathbf{P} = 0.5 \cdot \mathbf{e}_5 + 0.5 \cdot \mathbf{e}_6$ the optimal strategies show different behaviour ($\mathbf{e}_i$ is the i-th unit vector). The tendency to increase the sweep orders can be understood from a simple example: if the channel has no errors and a sweep of order 16 has revealed nothing, this implies that the FPAN coordinator does not have any of the beacon orders from 0 to 4. This is reflected in the updated state, and consequently the next sweep will have a larger order, because another sweep of order 16 or lower likely does not succeed.

5. Bounded targeted discovery

In this section we consider the case that only a limited time budget is available for discovery. When this budget has been exhausted without locating the FPAN coordinator, the device gives up. Such an approach makes sense when it is not certain that the required FPAN coordinator is present and one wants to limit the effort.

We first explain how the desired performance measures (discovery probability, average total listening costs) can be computed, then we provide some numerical results, including a validation of these results against simulation results.

5.1. Model description

We explain the approach for the special case of a strategy $\mathcal{S} = (s_1, s_2)$ with two stages, assuming that the FPAN coordinator to be discovered is indeed present and has chosen channel $F = f$. We are given an initial probability vector $\mathbf{P}$ for the beacon order of the foreign PAN. From Equation 2, the probability to find the foreign PAN after the first sweep is given by

$$\alpha(s_1, \mathbf{P}).$$

The number of channels below channel f is on average given by $G := \frac{F_{\max}-1}{2}$, and the average time the listener spends on these channels is therefore given by $G \cdot s_1$. With probability $\alpha(s_1, \mathbf{P})$ the foreign PAN is indeed found on channel f, and the cost incurred on this channel in case of success is denoted as $\beta(s_1, \mathbf{P})$, to be explained below. Afterwards the search stops. The foreign PAN is not found with probability $1 - \alpha(s_1, \mathbf{P})$, and in this case the listener listens for s_1 time slots on channel f, it then listens for further $G \cdot s_1$ time slots on the channels above f, then introduces a random waiting time of average duration $\frac{W_{\max}}{2}$, and finally starts the second sweep. In the second sweep first the time $G \cdot s_2$ is spent on channels below f, then the listener finds the foreign PAN on channel f with probability $\alpha(s_2, T_{s_1}(\mathbf{P}))$ and success-costs $\beta(s_2, T_{s_1}(\mathbf{P})$, otherwise it spends another $(G+1) \cdot s_2$ time slots for listening. Summarizing, the discovery probability is given by:

$$h(\mathcal{S}) = \alpha(s_1, \mathbf{P}) + (1 - \alpha(s_1, \mathbf{P})) \cdot \alpha(s_2, T_{s_1}(\mathbf{P})) \tag{5}$$

and the average total listening costs are given by:

$$\begin{aligned}
c(\mathcal{S}) = {} & G \cdot s_1 + \alpha(s_1, \mathbf{P}) \cdot \beta(s_1, \mathbf{P}) + (1 - \alpha(s_1, \mathbf{P})) \cdot \\
& \Big[(G+1) \cdot s_1 + \frac{W_{\max}}{2} + G \cdot s_2 + \alpha(s_2, T_{s_1}(\mathbf{P})) \cdot \beta(s_2, T_{s_1}(\mathbf{P})) \\
& + (1 - \alpha(s_2, T_{s_1}(\mathbf{P}))) \cdot (G+1) \cdot s_2\Big] \\
= {} & G \cdot s_1 + \alpha(s_1, \mathbf{P}) \cdot \beta(s_1, \mathbf{P}) + (1 - \alpha(s_1, \mathbf{P})) \cdot \left((G+1) \cdot s_1 + \frac{W_{\max}}{2}\right) \\
& + (1 - \alpha(s_1, \mathbf{P})) \cdot \Big[G \cdot s_2 + \alpha(s_2, T_{s_1}(\mathbf{P})) \cdot \beta(s_2, T_{s_1}(\mathbf{P})) \\
& + (1 - \alpha(s_2, T_{s_1}(\mathbf{P}))) \cdot (G+1) \cdot s_2\Big] \\
=: {} & c_I(s_1, \mathbf{P}) + (1 - \alpha(s_1, \mathbf{P})) \cdot c_F(s_2, T_{s_1}(\mathbf{P}))
\end{aligned} \tag{6}$$

where $c_I(s, \mathbf{P})$ can be regarded as the average cost incurred in an intermediate state (any state but the last) when using a sweep of order s and beacon order probability vector $\mathbf{P}$, whereas $c_F(s, \mathbf{P})$ are the average cost in the final state. The costs $\beta(s, \mathbf{P})$ are defined as:

$$\beta(s, \mathbf{P}) = \frac{1}{2} \sum_{b=0}^{B_{\max}} P_b \cdot \min\left\{s, 2^b\right\}$$

The rationale for this is as follows: when the sweep order s is larger than the beacon period 2^b, it takes (because of the assumption on the phase shift Φ) on average $2^b/2$ slots for discovery, whereas for $s \leq 2^b$ it takes on average $s/2$ slots. A number of comments are in order:

- From Equation 3 it can be concluded that the strategies (s_1, s_2) and (s_2, s_1) have the same discovery probability. By induction, this carries over to strategies of arbitrary length, i.e. all permutations of $(s_1, \ldots, s_k)$ have the same discovery probability.
- This model again uses the "memory" obtained from the evolution of the state probability vectors expressed by replacing $\mathbf{P}$ with $T_{s_1}(\mathbf{P})$.
- For the case of three states $\mathcal{S} = (s_1, s_2, s_3)$ we get in a similar fashion for the success probability:

$$\begin{aligned}
h(\mathcal{S}) &= \alpha(s_1, \mathbf{P}) + (1 - \alpha(s_1, \mathbf{P})) \cdot \Big(\alpha(s_2, T_{s_1}(\mathbf{P})) \\
&\quad + (1 - \alpha(s_2, T_{s_1}(\mathbf{P})) \cdot \alpha(s_3, T_{s_2}(T_{s_1}(\mathbf{P}))))\Big)
\end{aligned} \tag{7}$$

  and the average total listening costs:

$$\begin{aligned}
c(\mathcal{S}) = {} & c_I(s_1, \mathbf{P}) + (1 - \alpha(s_1, \mathbf{P})) \cdot \\
& \Big(c_I(s_2, T_{s_1}(\mathbf{P})) + (1 - \alpha(s_2, T_{s_1}(\mathbf{P}))) \cdot c_F(s_3, T_{s_2}(T_{s_1}(\mathbf{P})))\Big)
\end{aligned} \tag{8}$$

  It is straightforward to extend this to larger numbers of stages.

In the remainder of the paper we restrict to the cases of two and three states, since we believe that more states are not practical for relatively limited time budgets of, for example, $16 \cdot 100$ slots (approximately 24 seconds). With larger numbers of stages for a given cost budget too much time would be spent with actual channel switching or in the random waiting times at the end of a sweep.

5.2. Validation and numerical examples

We validate the analytical model against a simple simulation model, using the example of two-stage strategies (s_1, s_2). The simulation model assumes that the FPAN coordinator is indeed present and there are no packet losses. The FPAN picks one of the $F_{\max} = 16$ channels randomly according to a uniform distribution. In the first set of experiments we use different fixed beacon orders for the FPAN, namely $b \in \{6, 7, 8\}$, the listener uses two states with $n = 100$ slots per channel, split such that $s_1 = n - k$ and $s_2 = k$, for k varying from 1 to 99. For each parameter set the simulation has been repeated 3,000,000 times. The parameter $W_{\max}$ has always been set to 2^b in order to achieve perfect independence among different visits on the same channel. The results are shown in Figure 4. It can be seen that the simulation and analytical results for both the detection probability and the average total listening time show almost perfect agreement.

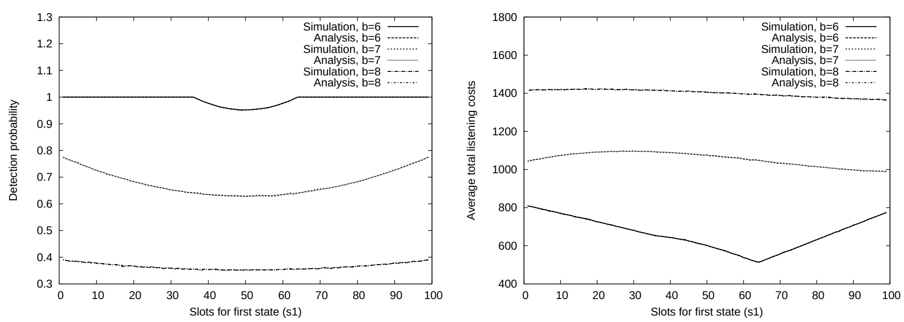

(a) Detection probability　　**(b)** Average total listening costs

Figure 4: Detection probability and average total listening costs (analysis and simulation) for all possible two-stage strategies with $s_1 = n - k$, $s_2 = k$, $n = 100$ and $k \in \{1, 2, \ldots, 99\}$ for different beacon orders $b \in \{6, 7, 8\}$ and $W_{\max} = b$. The lowest two curves in the left figure (which appear to be one curve because they overlap almost perfectly) represent the case $b = 8$, the middle two curves the case $b = 7$ and the remaining two curves the case $b = 6$. In the right figure the case $b = 8$ is represented by the upper two curves.

To further validate the analytical model, we have also performed experimental evaluations with an implementation on the Tmote Sky [17] mote platform using the TinyOS 2 [18] operating system and the ChipCon CC2420 transceiver

	Detection Probability	Avg. Costs
Model, strategy $\{2,5,6\}$	0.8878	24.977
Measurements, strategy $\{2,5,6\}$	0.8933	24.8
Model, strategy $\{3,5\}$	0.779	15.10
Measurements, strategy $\{3,5\}$	0.7789	15.33

Table 4: Comparison of model prediction with measurement results for two selected listening strategies, and $F_{\max} = 8$, $B_{\max} = 8$

[14] (see [19]). These measurements have been made in a setup where both the listener and the FPAN (represented by one single coordinator node) were very close to each other, so that almost no packet losses have been observed. For the measurements we have set $F_{\max} = 8$ and $B_{\max} = 8$. We have investigated two different strategies, namely $\mathcal{S} = \{4, 32, 64\}$ and furthermore $\mathcal{T} = \{8, 32\}$. For each of these strategies 10,000 independent repetitions of the experiment have been made. The experiments have been carried out at a weekend in our institute building to reduce the influence of external interference. The results are listed in Table 4. There is a very good agreement between model and experiments.

In the second validation experiment we have again considered two-stage strategies, but now the FPANs beacon order is random and drawn from a uniform distribution over $\{0, \ldots, 14\}$. We have used different fixed values for the $W_{\max}$ parameter. The simulations have been repeated 3,000,000 times for each parameter set. The results are shown in Figure 5. It can be observed that there are deviations between model and simulation results, but for the average total listening costs both model and simulation predict the cost-optimal strategy $\mathcal{S} = (16, 84)$. To explain the deviations, we consider one specific example where $W_{\max}$ is too small. For $W_{\max} = 0$, $F_{\max} = 2$ and $\mathcal{S} = (16, 16, 16, 16, \ldots)$ and the foreign PANs beacon order being $B = 5$, the simulation gives a success probability of 0.5, whereas the analytical model (which does not depend at all on $W_{\max}$) gives 1. In this example, on a single channel the FPANs beacons and the listening periods always have the same phase to each other – if the FPAN is not detected in the first sweep, it is never detected. The randomization is therefore strictly needed to ensure discovery. The $W_{\max}$ parameter determines how quickly the beacons and the listening periods get "out of synch".

We finally present results for the optimal strategies under different parameter sets. We have considered cost budgets of 100 and 200 slots on each channel, strategies with one, two or three stages, varying packet error rates r and uniform initial probabilities for the FPAN beacon order. The results have been obtained from the analytical model by exhaustively considering all possible two- and three-stage strategies. The results are shown in Table 5. It can be observed that for the considered parameter sets the optimal strategies do not depend on the packet error probability r. Secondly, the cost-optimal two- and three-stage-policies have significantly lower costs than the single-stage policy. Furthermore, similar to the case of unbounded costs, all cost-optimal strategies with at least two stages start with sweep orders that are powers of two, and, while not shown

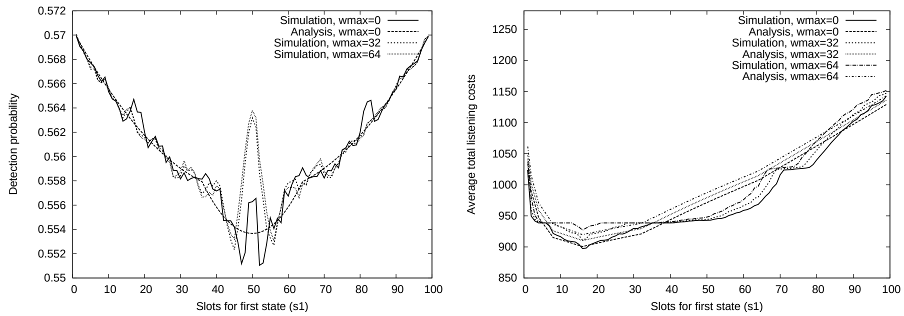

(a) Detection probability

(b) Average total listening costs

Figure 5: Detection probability and average total listening costs (analysis and simulation) for all possible two-sweep strategies with $s_1 = n - k$, $s_2 = k$, $n = 100$ and $k \in \{1, 2, \ldots, 99\}$ for beacon order drawn randomly with uniform distribution..

here, the success probabilities achieved with these optimal strategies deviate from the optimal success probability by less than one percent. Vice versa, as for example shown in Figure 5, the strategies that are optimal in terms of success probabilities can have average total listening costs that are *significantly* higher than the costs for the cost-optimal strategy. This is also true for the other scenarios considered in Table 5.

6. Further variations

In this section we consider two further variations of the discovery problem: unbounded untargeted discovery (i.e. discovering arbitrary other BSNs), and the usage of multiple listeners.

6.1. Discovering arbitrary PANs

The results presented in this section are motivated by opportunistic message relaying applications. In the scope of the ANGEL project[6] we have designed and implemented a protocol, the *critical message delivery protocol* (CMDP) [19], which provides the major building blocks to implement opportunistic message

[6]ANGEL - "Advanced Networked embedded platform as a Gateway to Enhance quality of Life" was a research project partially supported by the European Commission within the 6th Framework Programme (IST project 2005-IST-5-033506-STP), see `http://www.ist-angel-project.eu`. ANGEL aims at providing methods and tools for building networked embedded systems to monitor and improve the quality of life in common habitats, e.g., home, car and city environment. A key ingredient of ANGEL are body sensor networks based on IEEE 802.15.4 and ZigBee. These BSNs move around with their users and occasionally offload data to fixed or mobile gateways. Such a setup can for example be exploited for opportunistic message relaying applications.

Budget	Num. stages	Opt. succ. prob.	Opt. cost
PER $r = 0$			
Budget	**Num. stages**	**Opt. succ. prob.**	**Opt. cost**
100	1	$\mathcal{S} = (100)$, $h(\mathcal{S}) \approx 0.57$	$\mathcal{S} = (100)$, $c(\mathcal{S}) \approx 1133$
100	2	$\mathcal{S} = (99, 1)$, $h(\mathcal{S}) \approx 0.57$	$\mathcal{S} = (16, 84)$, $c(\mathcal{S}) \approx 910$
100	3	$\mathcal{S} = (98, 1, 1)$, $h(\mathcal{S}) \approx 0.57$	$\mathcal{S} = (8, 32, 60)$, $c(\mathcal{S}) \approx 894$
200	1	$\mathcal{S} = (200)$, $h(\mathcal{S}) \approx 0.636$	$\mathcal{S} = (200)$, $c(\mathcal{S}) \approx 2153$
200	2	$\mathcal{S} = (199, 1)$, $h(\mathcal{S}) \approx 0.636$	$\mathcal{S} = (32, 168)$, $c(\mathcal{S}) \approx 1613$
200	3	$\mathcal{S} = (198, 1, 1)$, $h(\mathcal{S}) \approx 0.636$	$\mathcal{S} = (8, 64, 128)$, $c(\mathcal{S}) \approx 1560$
PER $r = 0.1$			
Budget	**Num. stages**	**Opt. succ. prob.**	**Opt. cost**
100	1	$\mathcal{S} = (100)$, $h(\mathcal{S}) \approx 0.556$	$\mathcal{S} = (100)$, $c(\mathcal{S}) \approx 1144$
100	2	$\mathcal{S} = (99, 1)$, $h(\mathcal{S}) \approx 0.556$	$\mathcal{S} = (16, 84)$, $c(\mathcal{S}) \approx 930$
100	3	$\mathcal{S} = (98, 1, 1)$, $h(\mathcal{S}) \approx 0.556$	$\mathcal{S} = (8, 32, 60)$, $c(\mathcal{S}) \approx 913$
200	1	$\mathcal{S} = (200)$, $h(\mathcal{S}) \approx 0.62$	$\mathcal{S} = (200)$, $c(\mathcal{S}) \approx 2175$
200	2	$\mathcal{S} = (199, 1)$, $h(\mathcal{S}) \approx 0.62$	$\mathcal{S} = (32, 168)$, $c(\mathcal{S}) \approx 1653$
200	3	$\mathcal{S} = (198, 1, 1)$, $h(\mathcal{S}) \approx 0.623$	$\mathcal{S} = (8, 64, 128)$, $c(\mathcal{S}) \approx 1600$
PER $r = 0.2$			
Budget	**Num. stages**	**Opt. succ. prob.**	**Opt. cost**
100	1	$\mathcal{S} = (100)$, $h(\mathcal{S}) \approx 0.54$	$\mathcal{S} = (100)$, $c(\mathcal{S}) \approx 1156$
100	2	$\mathcal{S} = (99, 1)$, $h(\mathcal{S}) \approx 0.54$	$\mathcal{S} = (16, 84)$, $c(\mathcal{S}) \approx 952$
100	3	$\mathcal{S} = (98, 1, 1)$, $h(\mathcal{S}) \approx 0.54$	$\mathcal{S} = (8, 32, 60)$, $c(\mathcal{S}) \approx 936$
200	1	$\mathcal{S} = (200)$, $h(\mathcal{S}) \approx 0.608$	$\mathcal{S} = (200)$, $c(\mathcal{S}) \approx 2199$
200	2	$\mathcal{S} = (199, 1)$, $h(\mathcal{S}) \approx 0.608$	$\mathcal{S} = (32, 168)$, $c(\mathcal{S}) \approx 1697$
200	3	$\mathcal{S} = (198, 1, 1)$, $h(\mathcal{S}) \approx 0.608$	$\mathcal{S} = (8, 32, 60)$, $c(\mathcal{S}) \approx 1645$

Table 5: Optimal strategies for varying packet error rate, varying budget and varying number of stages. Fixed parameters: $W_{\max} = 32$, initial probability vector $\mathbf{P}$ chosen as uniform distribution.

relaying schemes [7], [8], [9]. In such schemes one BSN exploits the mobility of others to propagate messages to their destination. These building blocks include the discovery of foreign BSNs and the transmission of data packets into (discovered) foreign BSNs. One distinctive feature of the CMDP protocol is the ability to use multiple listener nodes for discovery and data delivery and to delegete the listening tasks to helper nodes selected by the BSN coordinator.

In this section we consider the discovery component in a mobile scenario. We consider one PAN (called the *listener*) that wishes to find other PANs. The listener uses a single node for listening, and this node perpetually performs a given sweep strategy $\mathcal{S} = \{s_1, s_2, \ldots, s_k\}$. The main performance measure of interest is the average time needed between the arrival of a new message at the listener and the first detection of a neighbor to which the message can be copied (*discovery time*). The listener searches all the time, until the first neighbor is discovered.

We have addressed this question by simulation. We first describe the simulation setup and then briefly present some results.

6.1.1. Simulation setup

A simplified version of the discovery component of the CMDP protocol has been implemented in the OMNet++ network simulator version 3.3[7] using the mobility framework version 2.0p3[8]. Our discovery protocol operates on top of an IEEE 802.15.4 MAC and PHY model. The mobility framework provides support for different radio channel and mobility models.

We have considered a square area of 3000×3000 m^2 in which the listening PAN plus a number $N = 50, 100, \ldots, 450$ of mobile FPANs are placed. The initial positions of all PANs (including the listener) are chosen randomly according to a uniform distribution. The listening PAN is stationary, all other N foreign PANs are mobile. They move independently of each other according to a variant of the random waypoint mobility model. The minimum and maximum speed of the foreign PANs has been set to 3 m/s and 5 m/s, respectively, with a zero pause time after reaching the destination.[9] Each foreign PAN initially picks a speed uniformly distributed between the minimum and maximum speed. The phase shifts of the FPANs are choosen uniformly out of their corresponding beacon intervals.[10]

The listener always starts at the first frequency channel. Each foreign PAN independently picks a frequency channel out of 16 available channels according

[7] `http://www.omnetpp.org`

[8] `http://mobility-fw.sourceforge.net/`

[9] The choice of a minimum speed removes a known problem of the random waypoint model, see [20], [21].

[10] These settings allow to start the simulation already under steady-state conditions. One exception is the distribution of the mobile speeds, which after some simulation time deviates from the initial uniform distribution, but is still close enough to the uniform distribution. We believe that our results are valid despite this minor inaccuracy. It would have required excessive simulation time for the speeds to converge to their stationary distribution.

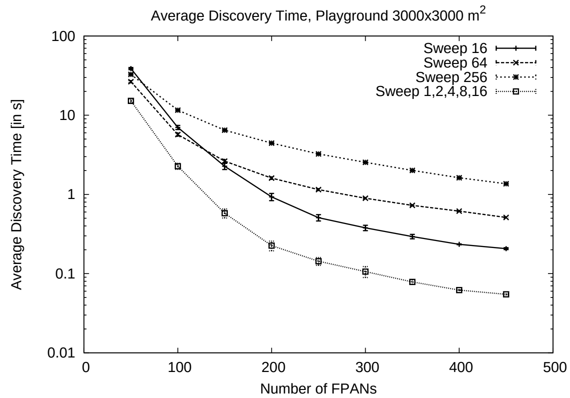

Figure 6: Average discovery time (in seconds) in mobile scenarios with varying number of mobile FPANs

to a uniform distribution, and furthermore picks independently a beacon order from 0 to 14, again according to a uniform distribution. The IEEE 802.15.4 model uses the default MAC parameters suggested in the standard. The radio model is a unit disc model, based on a log-distance channel model with a path loss exponent of two and a transmit power of 1mW for all stations.

Besides the number N of FPANs we have mainly varied the (perpetually executed) sweep strategy $\mathcal{S}$. We have used the single-sweep strategies (16), (32), ..., (512), and the multiple-sweep strategy $(1, 2, 4, 8, 16)$, the selection of which was inspired by the results reported in Section 4.2.

For each parameter set a number of 50,000 replications has been carried out. A replication is terminated when the listener discovers the first FPAN, and the time until detection as well as the beacon order of the discovered PAN have been recorded.

6.1.2. Results

In Figure 6 we show the average detection delay (in seconds) for varying number N of mobile FPANs (the figure includes confidence intervals for a confidence level of 95%). Furthermore, in Figure 7 we show the relative frequencies of the beacon orders of discovered FPANs for FPAN densities of $N = 50$ and $N = 300$, respectively (the curves for other values of N are similar). The following points are noteworthy:

- The multiple-sweep strategy $(1, 2, 4, 8, 16)$ has consistently the smallest average delay of all considered strategies. While due to the complexity of the scenario we cannot hope to prove any strategy to be optimal, this result suggests that multiple-sweep strategies using increasing powers of two (compare Table 2) are good candidates.

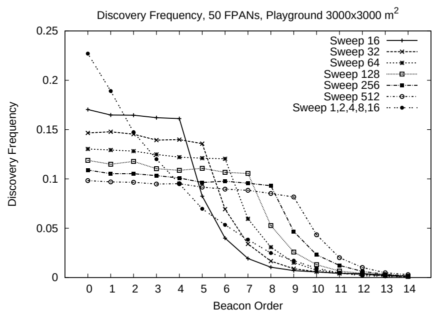

(a) 50 FPANs

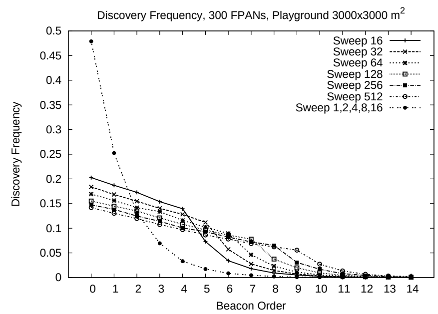

(b) 300 FPANs

Figure 7: Relative frequencies of beacon orders of discovered FPANs in a mobile scenario with 50 and 300 mobile FPANs, respectively.

- All strategies preferrably find mobile FPANs with small beacon orders (Figure 7). This results from the fact that FPANs with smaller beacon orders send more beacons per unit time than FPANs with larger beacon orders, and therefore have more chances to be detected. This is also beneficial for opportunistic message relaying applications, since FPANs with small beacon orders promise smaller delays in the further processing of the message.

- In Figure 7, all strategies consisting of a single sweep of order 2^a show the same behaviour: the relative frequencies first decay slowly from 0 to a, and then decay quickly, in a geometrically-looking fashion. Let us call the first part the "lower" part, the second part the "upper" part. For low FPAN densities one would expect a uniform distribution in the lower part, since with high probability only one mobile FPAN at a time is in the vicinity of the listener, and for this FPAN any beacon order is equiprobable. Beacon orders $\leq a$ are reliably detected by our channel model. When the number of mobile FPANs is increased, the deviation from the uniform distribution in the lower part increases as well. To understand this, consider the regime in which with high probability at most two FPANs are in the listeners range at any time. Suppose that these have different beacon orders $\leq a$. If in addition they have different frequencies, either one of them is discovered with the same probability, again giving rise to a uniform distribution. However, if they operate on the same frequency, the FPAN with the smaller beacon order is more likely detected, this way introducing a distortion to the uniform distribution. The magnitude of this distortion increases with the number of mobile FPANs, as it becomes more and more likely to have two or more FPANs sharing a channel.

6.2. Using multiple listeners

In [19] we have considered different options for the usage of multiple listeners. In one of them (the "overlapping scheme"), the listeners all follow the same listening strategy, but each listener is assigned a separate start channel. Instead of having all listeners start on channel one, the start channels are evenly spaced. We have investigated the benefits of this approach experimentally, utilizing our CMDP implementation on the Tmote Sky mote platform (see based on TinyOS 2 and the ChipCon CC2420 transceiver. Our experiments have shown that the search process can suffer from hardware varieties: two nominally identical transceivers (i.e. being the same product from the same vendor) placed at the same location can show significantly different receive sensitivities. We somewhat placatively call this "hardware diversity". With the overlapping strategy, the same channel is visited by multiple different listeners, providing a chance to compensate for hardware diversity. It is also intuitively clear that with multiple, geographically separated listeners visiting the same channel, spatial diversity [22] can be exploited, so that in fading channel scenarios the chances to discover a foreign PAN are increased. Finally, our experiments have also indicated (as could be expected) that with multiple listeners the average time until discovery can be reduced.

7. Related work

In summary, there is relatively few work addressing the (passive) discovery of IEEE 802.15.4-based BSNs [23] in more detail. To the best of our knowledge, there is no work which investigates the usage of multiple listening nodes for BSN discovery.

The temporary interconnection of ZigBee PANs (utilizing IEEE 802.15.4 as underlying technology) is described in [24]. ZigBee operates with the non-beacon-enabled mode of IEEE 802.15.4 and therefore active discovery can be used. The approach to find a PAN in radio range depends on which channels the PANs are operating. If both use the same channel, PANs should detect each other by receiving beacon frames with the PAN coordinator subfield set and not matching the stored PAN Id or coordinator address. In case two PANs are operating on different channels the use of the active and energy detection scan is proposed without suggesting any discovery strategy.

In [25] the authors describe an asynchronous neighbor discovery and rendezvous protocol called Disco (not related to IEEE 802.15.4). The protocol allows nodes to run at low duty-cycle, but they are still able to discover each other without any synchronization. Each node picks a pair of prime numbers such that the sum of the reciprocals is equal to the desired duty cycle. Nodes turn their radio on and start to transmit and receive whenever a node-local counter is divisible by either of the primes. This adaption of the Chinese Remainder Theorem ensures the discovery in bounded time.

An analytical framework based on a Markov Model for finding a tradeoff between energy efficiency and discovery timeliness for neighborhood detection in self-organizing ad hoc and sensor networks is developed in [26] [27]. Nodes searching for others nodes in reach perform a set of procedures which is called *Hunting Process*. The *Hunting Process* consists of two modes: the Inquiry and the Inquiry Scan mode. In the Inquiry mode the node broadcasts beacon messages in order to enable the detection by other nodes that are in the Inquiry Scan mode. However, the authors put their work into the scope of ad hoc networks (see also [28]) and do not specifically consider the IEEE 802.15.4 WPAN.

There has also been some work on assessing the discovery times for other PAN technologies. Specifically, for Bluetooth in [29], [30] an analytical assessment of the discovery times is done and guidelines for configuring Bluetooth to enable quicker discovery are derived.

8. Conclusions

In this paper we have presented a detailed assessment of passive discovery of IEEE 802.15.4 BSNs for a specific class of listening strategies. For the case of targeted discovery we have given models allowing to derive optimal search strategies. For bounded targeted discovery we have observed that the strategies optimizing the success probability and the listening costs are not the same. The cost-optimal strategy often has success probabilities that are not much smaller than the one of the success-optimal strategies, but vice versa the success-optimal strategy is significantly more costly than the cost-optimal strategy. In the unbounded untargeted scenario it is preferrable to stay for longer time on one channel. When the foreign BSNs uses higher beacon orders, they are only rarely discovered and mostly “hidden”. To support opportunistic message relaying applications, it is therefore beneficial to avoid the large beacon orders.

There are many opportunities for future research. The first is the design of active discovery schemes that are interoperable with the existing IEEE 802.15.4 standard. One theoretically interesting challenge is a precise characterization of the *distribution* of the discovery time for a given strategy $\mathcal{S}$, not only the computation of its average. In another direction, the design of energy-efficient and reliable listening strategies to detect *any* other network (not only IEEE 802.15.4 networks but also WiFi networks) and to assess the utilization of a channel as a preparation for channel selection is a very interesting research topic. Finally, it will probably be necessary in many applications to equip the coordinator (or the nodes) of a searching PAN with a pre-defined set of strategies for different situations. A key challenge is then of course the design of this fixed set of strategies.

References

[1] A. Milenkovic, C. Otto, E. Jovanov, Wireless sensor networks for personal health monitoring: Issues and an implementation, Computer CommunicationsTo appear.

[2] A. Natarajan, M. Motani, B. de Silva, K.-K. Yap, K. C. Chua, Investigating network architectures for body sensor networks, in: Proc. 1st ACM SIGMOBILE international workshop on Systems and networking support for healthcare and assisted living environments (HealthNet), 2007.

[3] H. Karl, A. Willig, Protocols and Architectures for Wireless Sensor Networks, John Wiley & Sons, Chichester, 2005.

[4] LAN/MAN Standards Committee of the IEEE Computer Society, IEEE Standard for Information technology – Telecommunications and information exchange between systems – Local and metropolitan area networks – Specific requirements – Part 15.4: Wireless Medium Access Control (MAC) and Physical Layer (PHY) Specifications for Low Rate Wireless Personal Area Networks (LR-WPANs), revision of 2006 (Sep. 2006).

[5] R. C. Shah, L. Nachman, C. yih Wan, On the performance of bluetooth and ieee 802.15.4 radios in a body area network, in: Proc. BodyNets'08, Tempe, USA, 2008.

[6] E. Miluzzo, X. Zheng, K. Fodor, A. T. Campbell, Radio characterization of 802.15.4 and its impact on the design of mobile sensor networks, in: Proc. Wireless Sensor Networks, Fifth European Workshop (EWSN 2008), Bologna, Italy, 2008.

[7] L. Pelusi, A. Passarella, M. Conti, Opportunistic networking: Data forwarding in disconnected mobile ad hoc networks, IEEE Communications Magazine 44 (11) (2006) 134 – 141.

[8] M. Conti, S. Giordano, Multihop ad hoc networking: The theory, IEEE Communications Magazine 45 (4) (2007) 78 – 86.

[9] M. Conti, S. Giordano, Multihop ad hoc networking: The reality, IEEE Communications Magazine 45 (4) (2007) 88 – 95.

[10] R. C. Shah, S. Roy, S. Jain, W. Brunette, Data mules: Modeling a three-tier architecture for sparse sensor networks, in: Proc. IEEE Workshop on Sensor Network Protocols and Applications (SNPA), 2003.

[11] P. Juang, H. Oki, Y. Wang, M. Martonosi, L.-S. Peh, D. Rubenstein, Energy-efficient computing for wildlife tracking: Design tradeoffs and early experiences with zebranet, in: Proc. 10th International Conference on Architectural Support for Programming Languages and Operating Systems, San Jose, CA, 2002.

[12] M. Grossglauser, D. N. C. Tse, Mobility increases the capacity of ad hoc wireless networks, IEEE/ACM Transactions on Networking 10 (4) (2002) 477–486.

[13] Z. J. Haas, T. Small, A New Networking Model for Biological Applications of Ad Hoc Sensor Networks, IEEE/ACM Transactions on Networking 14 (1) (2006) 27–40.

[14] Chipcon, 2.4 GHz IEEE 802.15.4 / ZigBee-ready RF Transceiver, Chipcon Products from Texas Instruments (2004).

[15] S. M. Ross, Introduction to Stochastic Dynamic Programming, Academic Press, San Diego, 1983.

[16] D. P. Bertsekas, Dynamic Programming and Optimal Control – Volume 1, 3rd Edition, Athena Scientific, Belmont, Massachusetts, 2005.

[17] J. Polastre, R. Szewczyk, D. Culler, Telos: Enabling ultra-low power wireless research, in: Proc. Information Processing in Sensor Networks, Special track on Platform Tools and Design Methods for Network Embedded Sensors (IPSN/SPOTS), Los Angeles, CA, 2005.

[18] P. Levis, D. Gay, V. Handziski, J.-H.Hauer, B.Greenstein, M.Turon, J.Hui, K.Klues, C.Sharp, R.Szewczyk, J.Polastre, P.Buonadonna, L.Nachman, G.Tolle, D.Culler, A.Wolisz, T2: A second generation os for embedded sensor networks, Technical Report TKN-05-007, Telecommunication Networks Group, Technische Universität Berlin (Nov. 2005).

[19] N. Karowski, A. Willig, J. Hauer, Passive discovery schemes for opportunistic message relaying schemes based on ieee 802.15.4, TKN Technical Report Series TKN-08-008, Telecommunication Networks Group, Technical University Berlin (Aug. 2008).

[20] C. Bettstetter, G. Resta, P. Santi, The node distribution of the random waypoint mobility model for wireless ad hoc networks, IEEE Transactions on Mobile Computing 2 (3) (2003) 257–269.

[21] J. Yoon, M. Liu, B. Noble, Random waypoint considered harmful, in: Proc. IEEE Infocom, San Francisco, California, 2003.

[22] S. N. Diggavi, N. Al-Dhahir, A. Stamoulis, A. R. Calderbank, Great Expectations: The Value of Spatial Diversity in Wireless Networks, Proceedings of the IEEE 92 (2) (2004) 219–270.

[23] J. Misic, V. B. Misic, Wireless Personal Area Networks with IEEE 802.15.4, John Wiley & Sons, Chichester, UK, 2008.

[24] S. Jung, A. Chang, M. Gerla, Temporary interconnection of zigbee personal area networks (pan), in: Proc. Fourth Annual International Conference on Mobile and Ubiquitous Systems: Networking & Services (MobiQuitous 2007), Philadelphia, USA, 2007.

[25] P. Dutta, D. Culler, Practical asynchronous neighbor discovery and rendezvous for mobile sensing applications, in: SenSys '08: Proceedings of the 6th ACM conference on Embedded network sensor systems, ACM, New York, NY, USA, 2008, pp. 71–84.

[26] L. Gallucio, G. Morabito, S. Palazzo, Analytical evaluation of a tradeoff between energy efficiency and responsiveness of neighbour discovery in self-organizing ad hoc networks, IEEE Journal on Selected Areas in Communications 22 (7) (2004) 1167–1182.

[27] L. Gallucio, A. Leonardi, G. Morabito, S. Palazzo, Tradeoff between Energy-Efficiency and Timeliness of Neighbor Discovery in Self-Organizing Ad Hoc and Sensor Networks, in: Proc. 38th Annual Hawaii International Conference on System Sciences (HICSS '05), Hawaii, USA, 2005.

[28] S. A. Borbash, M. J. McGlynn, Birthday protocols for low energy deployment and flexible neighbour discovery in ad hoc wireless networks, in: Proc. 2nc ACM Intl. Symposium on Mobile Ad Hoc Networking and Computing (MobiHoc) 2001, Long Beach, CA, 2001.

[29] T. Salonidis, P. Bhagwat, L. Tassiulas, R. LaMaire, Distributed Topology Construction of Bluetooth Personal Area Networks, in: Proc. IEEE Infocom 2001, Anchorage, Alaska, 2001.

[30] T. Salonidis, P. Bhagwat, L. Tassiulas, R. LaMaire, Distributed topology construction of bluetooth personal area networks, IEEE Journal on Selected Areas in Communications 23 (3) (2005) 633–643.