



Telecommunication
Networks Group

Technische Universität Berlin
Telecommunication Networks Group

Proofs and Performance Evaluation of Greedy Multi-Channel Neighbor Discovery Approaches

Niels Karowski, Konstantin Miller

karowski@tkn.tu-berlin.de, konstantin.miller@ieee.org

Berlin, June 2018

TKN Technical Report TKN-18-0003

TKN Technical Reports Series
Editor: Prof. Dr.-Ing. Adam Wolisz

Abstract

The accelerating penetration of physical environments by objects with information processing and wireless communication capabilities requires approaches to find potential communication partners and discover services. In the present work, we focus on passive discovery approaches in multi-channel wireless networks based on overhearing periodic beacon transmissions of neighboring devices which are otherwise agnostic to the discovery process. We propose a family of low-complexity algorithms that generate listening schedules guaranteed to discover all neighbors. The presented approaches simultaneously depending on the beacon periods optimize the worst case discovery time, the mean discovery time, and the mean number of neighbors discovered until any arbitrary in time. The presented algorithms are fully compatible with technologies such as IEEE 802.11 and IEEE 802.15.4. Complementing the proposed low-complexity algorithms, we formulate the problem of computing discovery schedules that minimize the mean discovery time for arbitrary beacon periods as an integer linear problem. We study the performance of the proposed approaches analytically, by means of numerical experiments, and by extensively simulating them under realistic conditions. We observe that the generated listening schedules significantly – by up to factor 4 for the mean discovery time, and by up to 300% for the mean number of neighbors discovered until each point in time – outperform the Passive Scan, a discovery approach defined in the IEEE 802.15.4 standard. Based on the gained insights, we discuss how the selection of the beacon periods influences the efficiency of the discovery process, and provide recommendations for the design of systems and protocols.

1 Introduction

We are currently observing a rapid augmentation of physical objects surrounding us with information processing and wireless communication capabilities. It is estimated that by 2020 25 [1] up to 50 [2] billion objects will be connected to the Internet. This development is leading us to a new era of computing. The resulting network of “smart” objects that interact with each other and exchange information without a direct human intervention, the so-called Internet of Things (IoT), will serve as a foundation for novel applications in a wide range of domains.

In order to discover services of interest devices will need to detect other entities within communication range that are able to use common communication technology—the so called neighbors.

Neighbor discovery can be done in two fundamentally different ways. For an *active* discovery, the discoverer broadcasts probe requests that must be answered by the neighbors. An active discovery is fast but has the drawback that all neighbors have to consume energy by (continuously) listening to potential inquiries even though they might only be interested in being detected but not in discovering their own neighborhood. In contrast, *passive* schemes perform the discovery by overhearing beaconing messages that are periodically broadcasted (with a specific Beacon Period (BP)) by neighbors interested in being discovered. The beaconing neighbors themselves are hereby agnostic to the discovery process. Let us emphasize that periodic beaconing is already used in many widely deployed technologies such as IEEE 802.11 [3] and IEEE 802.15.4 [4]. In order to be compatible with current state-of-the-art technologies, such as IEEE 802.11 and IEEE 802.15.4, neighbor discovery must support multi-channel environments. Finally, we assume lack of time synchronization among the devices involved in the discovery process.

A frequently adopted objective for the design of discovery approaches is the minimization of the Worst-Case Discovery Time (WDT)—time required to detect all *potential* neighbors. A complete discovery is desirable, e.g., in order to avoid interference with neighbors when establishing a new network. In addition, minimizing the WDT has the advantage of implicitly minimizing the consumed energy. Other applications are interested in the maximization of the number of discoveries until a given point in time, which we call the Number of Discoveries over Time (NDoT), as, e.g., in the case of identifying potential forwarders in Delay Tolerant Networks. Yet other applications benefit from discovering the individual neighbors as early as possible, e.g., emergency services. Their objective is thus the minimization of the Mean Discovery Time (MDT). Since many devices in IoT environments will be battery powered, and will have limited computational resources, neighbor discovery should be performed in an energy-efficient way with low to moderate computational requirements.

The publication [5] associated with this report presents several novel contributions providing simple and efficient discovery algorithms applicable under realistic conditions:

We provide for the first time a full characterization of the class of listening schedules that are guaranteed to discover all neighbors (we call such schedules *complete*), and that pointwise maximize the Cumulative Distribution Function (CDF) of the discovery times. The latter feature implies that they optimize all three mentioned performance metrics *simultaneously*: WDT, MDT, and the NDoT. We call these schedules *recursive*, due to their specific structure.

Our second, practically most relevant, contribution consists of several approaches to construct listening schedules that, under certain assumptions, are recursive (and thus inherit the corresponding optimality properties). We define a family of low-complexity algorithms that we call GREEDY, due to their operation mode [6]. Further, we define an algorithm called CHAN TRAIN which is an extension of the GREEDY family, aiming at a reduction of the number of channel switches.

In general, the performance of discovery algorithms is strongly dependent on the allowed set of BP's, i.e. periods with which beacons are transmitted. To this point - incompatible with the state-of-the-art wireless protocols - assumptions about the beacon transmission patterns have been made so far in the literature. In this paper, we consider on one hand the most general case \mathbb{F}_1 , a family containing all possible BP sets, but introduce in addition also two other practically important families of BP sets— \mathbb{F}_2 and \mathbb{F}_3 . They include all BP sets supported by IEEE 802.15.4 and a large part of the BI sets supported by IEEE 802.11—two widely adopted standards for wireless communication.

We prove that for BP sets from \mathbb{F}_3 the listening schedules computed by GREEDY and CHAN TRAIN are recursive (and thus complete, and optimal w.r.t. the three targeted performance metrics). Moreover, for BP sets from \mathbb{F}_2 the computed schedules are complete and WDT-optimal, while they are close-to-optimal w.r.t. the MDT. Finally, we show that even for the most general case of \mathbb{F}_1 the computed schedule are complete, close-to-optimal w.r.t. the MDT, and still within 30% of the optimum for the WDT, while this gap decreases for an increasing number of channels.

Our third contribution demonstrates that even for arbitrary BP sets from \mathbb{F}_1 complete and MDT-optimal schedules are achievable, albeit only by solving an Integer Linear Program (ILP). We prove that computed schedules are also WDT-optimal for BP sets from \mathbb{F}_2 and NDoT-optimal for BP sets from \mathbb{F}_3 . This approach is attractive due to the broad range of supported BP sets. However, it has a high computational complexity and memory consumption, restricting its usage to offline computations, and to scenarios with a moderate number of channels and size of the used BP's.

As additional contribution, we define an algorithm called OPT_{B2} that computes recursive schedules for scenarios, in which the cardinality of the BP set is restricted to two entries.

Unfortunately performing of such discovery in real multi-channel environments suffers under an implementation impact: non-negligible deaf periods occur during the execution of a channel switch resulting in potentially missing some beacons transmitted during such deaf periods. Due to this effect even algorithms provably generating complete schedules will, in reality, miss some neighbors - the percentage of missed neighbors can reasonably be expected to increase with the increase of the number of channel switches required by a given algorithm. In order to quantify this impact, we perform simulations using a realistic wireless model and device behavior expressing the results in the form of an additional performance metric — the success rate, which is the fraction of neighbors discovered under this realistic conditions by any algorithm under consideration. Using this additional performance metric we have derived our next contribution: We suggest two instances of the GREEDY family of algorithms designed to reduce the number of channel switches and perform their simulative performance evaluation w.r.t achievable fraction of discovered neighbors.

In all evaluations, in addition to comparing the performance with the optimum, we perform a comparison against Passive Scan (PSV), a discovery scheme defined by the IEEE 802.15.4 standard. We observe that GREEDY algorithms significantly (by up to several hundreds percent) outperform PSV w.r.t. the MDT and the NDoT in all studied scenarios.

As our final contribution we discuss the strong impact the structure of allowed BP sets has on the performance of discovery approaches, and provide recommendations for a BP selection that supports efficient neighbor discovery. These recommendations may be useful, on the one hand, for the development of novel wireless communication based technologies that use periodic beaconing messages for management or synchronization purposes, and, on the other hand, for the BP selection for existing technologies that support a wide range of BP's, such as IEEE 802.11.

This report is a supplement to [5] and is structured as follows. In Section 2 additional examples of the properties of the developed discovery strategies are provided. Section 3 provides a proof on the

optimal duration of a complete discovery. Properties of recursive schedules regarding the optimality relating to the analyzed performance metrics are described in Section 4. A performance analysis for the GREEDY algorithms is provided in Section 5, for the CHAN TRAIN algorithm in Section 6 and for the OPT_{B2} algorithm in Section 7. Section 8 provides a proof on the WDT of MDT-optimal schedules. The relation between schedules being NDoT-optimal and WDT-optimal as well as MDT-optimal is discussed in Section 9. The computational complexity and memory requirements of a possible GREEDY implementation is presented in Section 10. Finally, Section 11 provides additional simulation results.

2 Examples

In the following figures gray squares represent the scanned channels, while the numbers in each square represent the expected value of the fraction of neighbors that can be discovered by scanning a particular channel during a particular time slot. In addition to the assumptions on the uniform distribution of channels and beacon offsets described in Section III in [5], in the following examples we also assume a uniform distribution of BP's. That is, the configuration probabilities are assumed to be $P_\kappa = \frac{1}{b_\kappa |B| |C|}$, for each $\kappa \in K_{BC}$. Finally, we remark that while the presented examples consider very small scenarios, the illustrated performance gaps may become arbitrarily large if larger scenarios are considered.

Example 1 (A recursive schedule does not always exist). *Consider the setting depicted in Figure 1, with $B = \{1, 2, 3\} \notin \mathbb{F}_2$ and $|C| = 2$. Vertical dashed lines mark time instants until which all configurations (c, b_i, δ) for a BP b_i have to be discovered on each channel $c \in C$ in order for the schedule to be recursive. Consequently, gray boxes indicate channels that have to be scanned during the first 4 time slots (uniquely determined up to swapping the channels). Observe that a recursive listening schedule for this example does not exist due to the fact that it would have to discover all remaining neighbor configurations with $b = 3$ during the two time slots 5 and 6. However, the number of remaining configurations is 3: $(0, 3, 0)$, $(1, 3, 1)$, and $(1, 3, 2)$.*




Figure 1: Illustration for Example 1. With $B = \{1, 2, 3\} \notin \mathbb{F}_2$ and $|C| = 2$, a recursive schedule does not exist.

Example 2 (GREEDY is not optimal for arbitrary BP sets). *For arbitrary BP sets, a GREEDY algorithm does not necessarily generate MDT-optimal or WDT-optimal schedules. Consider a setting with $B = \{1, 2, 3\} \in \mathbb{F}_1 \setminus \mathbb{F}_2$ and $|C| = 3$. Figure 2a shows a GREEDY listening schedule, which is neither WDT-optimal nor MDT-optimal, while Figure 2b depicts a schedule, which is both WDT-optimal and MDT-optimal but not GREEDY. In particular, optimal value for the WDT in this example*

constitutes 9 time slots, while optimal value for MDT is $2.6\bar{1}$ time slots, in contrast to $2.7\bar{2}$ time slots achieved by the schedule computed by GREEDY. Another example using $B = \{2, 3, 4, 6, 12\} \in \mathbb{F}_2 \setminus \mathbb{F}_3$ and $|C| = 2$ is depicted in Figure 3. The listening schedule generated by GREEDY is WDT-optimal but only achieves a MDT of 5.3 time slots in comparison with the optimal value of 5.1 time slots.

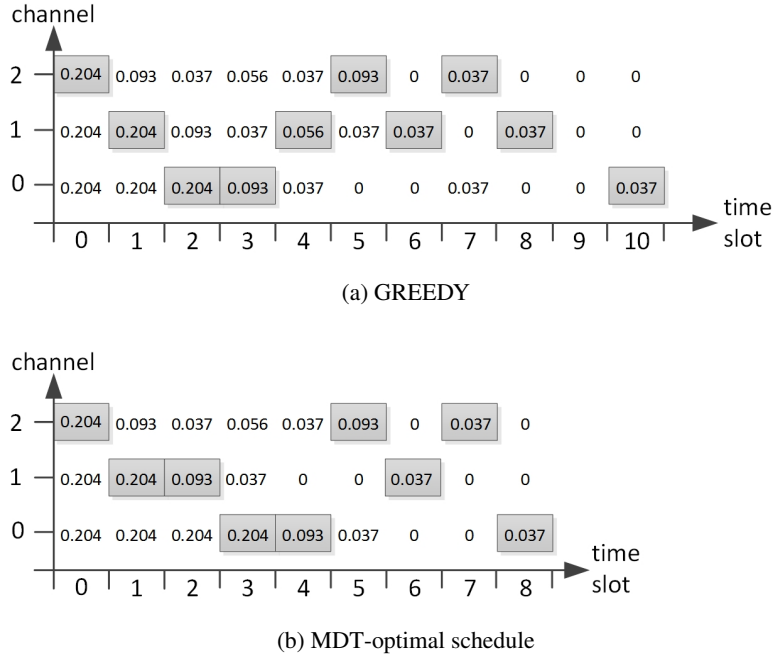



Figure 2: Illustration for Example 2 showing a schedule generated by GREEDY for $B = \{1, 2, 3\} \notin \mathbb{F}_2$ and $|C| = 3$ that is neither MDT-optimal nor WDT-optimal.

Example 3 (CHAN TRAIN is not GREEDY for general BP sets). For a BP set $B \notin \mathbb{F}_3$ Proposition 8 is no longer true. Figure 4 shows an example setting with $B = \{1, 2, 3, 6\} \in \mathbb{F}_2 \setminus \mathbb{F}_3$ and $|C| = 3$, in which this is the case.


Example 4 (MDTOPT is in general not WDT-optimal). Listening schedules generated by MDTOPT might not be WDT-optimal. Figure 5 depicts an example using $B = \{1, 2, 4, 5\}$ and $|C| = 2$, for which a MDT-optimal listening schedule cannot be constructed within $\max(B)|C|$ time slots. Figure 5a shows a MDT-optimal listening schedule with the additional constraint of using at most $\max(B)|C|$ time slots. Observe that its MDT is 2.875 time slots as compared to the optimal value of 2.75 time slots of the schedule shown in Figure 5b.

3 Optimal WDT

In this section, we prove that the optimum duration of a complete discovery is $\max(B)|C|$ time slots for arbitrary BP sets and channel sets. This allows us to define a WDT-optimal listening schedule as a schedule which requires $\max(B)|C|$ for a complete discovery.




(a) GREEDY




(b) MDT-optimal schedule

Figure 3: Illustration for Example 2 showing a schedule generated by GREEDY for $B = \{2, 3, 4, 6, 12\} \in \mathbb{F}_2 \setminus \mathbb{F}_3$ and $|C| = 2$ that is WDT-optimal but not MDT-optimal.



(a) CHAN TRAIN



(b) GREEDY

Figure 4: Illustration for Example 3 showing the non-greedy behavior of the CHAN TRAIN strategy for the BI set $B = \{1, 2, 3, 6\} \in \mathbb{F}_2 \setminus \mathbb{F}_3$ and $|C| = 3$.

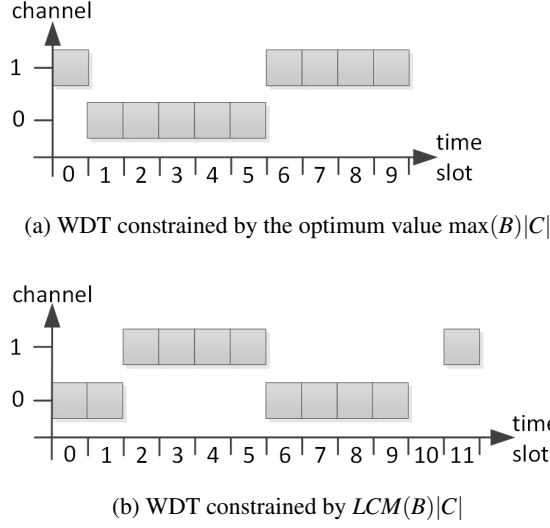


Figure 5: Illustration for Example 4, with BP set $B = \{1, 2, 4, 5\} \in \mathbb{F}_1 \setminus \mathbb{F}_2$ and $|C| = 2$. In this example, no schedule exists which is both MDT-optimal and WDT-optimal.

Proposition 1 (WDT-optimality). *For an arbitrary set of BP's B , and a set of channels C , the optimum duration of a complete discovery is $\max(B)|C|$ time slots.*

Proof. Observe that the complete discovery can always be performed within $\max(B)|C|$ time slots by listening on each channel for $\max(B)$ consecutive time slots. More precisely, channel c_j , $j \in \{0, \dots, |C| - 1\}$ is scanned during the time slots $\{jb_{\max}, \dots, (j+1)b_{\max} - 1\}$, where $b_{\max} = \max(B)$. Thus, configuration (c_j, b, δ) , $b \in B$, $\delta \in \{0, \dots, b - 1\}$, is discovered during the time slot $t = jb_{\max} + \delta$. Consequently, all configurations from K_{BC} are discovered during the time slots $\{0, \dots, b_{\max}|C| - 1\}$.

To see that a complete discovery cannot be performed faster than $\max(B)|C|$, consider the configurations $\{(c, b_{\max}, b_{\max} - 1) \mid c \in C\} \subseteq K_{BC}$. To discover them, it is necessary to scan at least the time slots $\{ib_{\max} + b_{\max} - 1\}_{i \in \{0, \dots, |C| - 1\}}$, proving the claim. \square

Consequently, we define a WDT-optimal schedule as follows.

Definition 2 (WDT-optimal schedule). *We call a listening schedule for a BP set B and a channel set C WDT-optimal if it discovers all neighbors within $\max(B)|C|$ time slots.*

Note that WDT-optimal schedules do not have idle time slots as stated in the following proposition.

Proposition 3 (WDT-optimality implies no idle slots). *WDT-optimal schedules have no idle time slots.*

Proof. Assume that a time slot $t \in \{0, \dots, b_{\max}|C| - 1\}$, where $b_{\max} = \max(B)$, B is the set of BP's, and C is the set of channels, is idle. Then, since the schedule is WDT-optimal, there remain $|C| - 1$ time slots $\{ib_{\max} + t \bmod b_{\max} \mid i \in \{0, \dots, |C| - 1\}\} \setminus \{t\}$, during which a configuration from the set $\{(c, b_{\max}, t \bmod b_{\max}) \mid c \in C\}$ can be discovered. Since the latter set of configurations has size $|C|$, we obtain a contradiction. \square

4 Recursive Schedules

Recursive schedules have a number of compelling properties that are stated in the following proposition.

Proposition 4 (Properties of recursive schedules). *A recursive schedule for a BP set $B \subset \mathbb{N}^+$ is complete, WDT-optimal, MDT-optimal, and NDoT-optimal.*

Proof. Note that recursive schedules are complete and WDT-optimal by definition (see Definition 1 in [5]).

Further, note that a schedule \mathcal{L} that maximizes the NDoT is also MDT-optimal since it pointwise maximizes the CDF of discovery times

$$\sum_{\kappa \in K_{BC}: \mathcal{T}_\kappa(\mathcal{L}) \leq t} P_\kappa \geq \sum_{\kappa \in K_{BC}: \mathcal{T}_\kappa(\mathcal{L}') \leq t} P_\kappa$$

for all $\mathcal{L}' \subset C \times \mathbb{N}, t \in \mathbb{N}$.

Consequently, it is sufficient to show NDoT-optimality. To prove it, we show that for a recursive schedule \mathcal{L} , each scan $(c, t) \in \mathcal{L}$ with $t \in \{0, \dots, b|C| - 1\}$ results in the discovery of the configuration $(c, b, t \bmod b)$. This implies that for each $b \in B$, the number of configurations detected until each time slot is maximized since it is not possible to discover more than one configuration with a certain BP $b \in B$ per time slot. Due to the assumption that the channels and the beacon offsets for each BP are uniformly distributed, it also implies that the sum of the discovery probabilities is maximized in each time slot, proving the claim.

Observe that for a $t \in \{0, \dots, b|C| - 1\}$, there are $|C|$ configurations with a BP b and an offset $t \bmod b$, namely the configurations $\{(c, b, t \bmod b) \mid c \in C\}$. At the same time, within the first $b|C|$ time slots there are exactly $|C|$ time slots where these configurations can be discovered, namely the time slots $\{ib + t \bmod b \mid i \in \{0, \dots, |C| - 1\}\}$. Consequently, since a recursive schedule per definition detects all configurations with the BP b during the first $b|C|$ time slots, it has to scan a different channel during each of the time slots $\{ib + t \bmod b \mid i \in \{0, \dots, |C| - 1\}\}$. It is thus detecting one configuration with the BP b and offset $t \bmod b$ in each of these time slots. \square

Unfortunately, not each scenario admits a recursive schedule, see Example 1 in Section 2. However, for BP sets from \mathbb{F}_3 recursive schedules always exist, as shown by the following proposition.

Proposition 5. *For a BP set $B \in \mathbb{F}_3$ a recursive schedule always exists.*

Proof. We prove the claim by induction. First, we show that a recursive schedule exists for a BP set containing just one element. Then, we show that a recursive schedule for a BP set $B' \in \mathbb{F}_3$ can be extended to a recursive schedule for a BP set $B'' \in \mathbb{F}_3$, with $B' \subset B''$, and $|B''| = |B'| + 1$.

To see that a recursive schedule for a BP set $B = \{b\}$ exists, consider the schedule that scans channel c_j , $j \in \{0, \dots, |C| - 1\}$, during the time slots $\{jb, \dots, (j+1)b - 1\}$. This schedule discovers all configurations during the time slots $\{0, \dots, |C|b - 1\}$, and is thus recursive.

Now, assume that we are given a recursive schedule for a BP set B' . We show that it can be extended to a recursive schedule for B'' . We denote $b'_{\max} = \max(B')$ and $b''_{\max} = \max(B'')$. W.l.o.g. we assume $b'' > b'$. Observe that $B'' \in \mathbb{F}_3$ implies $b'' = \alpha b'$, for an $\alpha \in \mathbb{N}_+$.

Observe that any complete and WDT-optimal schedule for a BP set B can be written in form of a matrix $A \in \mathbb{N}^{|C| \times b_{\max}}$, $b_{\max} = \max(B)$, where each element $A_{c\delta}$ indicates that channel c is scanned

during the time slot $A_{c\delta}b_{\max} + \delta$, and thus the configuration (c, b_{\max}, δ) is discovered during this time slot. For a matrix A to be a valid representation of a complete schedule, it must hold $A_{c\delta} \neq A_{c'\delta}$, for $c \neq c'$, $\delta \in \{0, \dots, b_{\max} - 1\}$, since otherwise two different channels would have to be scanned during the same time slot. Furthermore, it must hold $A_{c\delta} \in \{0, \dots, |C| - 1\}$ since a WDT-optimal schedule only uses time slots $t \in \{0, \dots, b_{\max}|C| - 1\}$.

Assume that the given recursive schedule for the BP set B' is represented by the matrix $A' \in \mathbb{N}^{|C| \times b'_{\max}}$. We need to show that A' can be transformed to a new matrix $A'' \in \mathbb{N}^{|C| \times b''_{\max}}$, such that A'' is a valid representation of a complete and WDT-optimal schedule, and such that the channels scanned during the time slots $t \in \{0, \dots, b'_{\max}|C| - 1\}$ by the schedule A' are also scanned during the same time slots by the schedule A'' .

We claim that such a transformation is obtained by mapping each element $A'_{c\delta'}$ to an element $A''_{c\delta''}$ as follows:

$$\begin{cases} \delta'' &= (A'_{c\delta'}b'_{\max} + \delta') \bmod b''_{\max} \\ A''_{c\delta''} &= \frac{(A'_{c\delta'}b'_{\max} + \delta') - \delta''}{b''_{\max}} \end{cases}.$$

Observe that since the numerator is an integer multiple of the denominator, the expression on the right-hand side evaluates to an integer.

This mapping retains the channels scanned by A' since

$$\begin{aligned} A''_{c\delta''}b''_{\max} + \delta'' &= (A'_{c\delta'}b'_{\max} + \delta') - \delta'' + \delta'' \\ &= A'_{c\delta'}b'_{\max} + \delta'. \end{aligned}$$

Moreover, this mapping is injective, such that no two elements from A' are mapped to the same element of A'' . This is due to the fact that $\delta'' = (A'_{c\delta'}b'_{\max} + \delta') \bmod b''_{\max}$ is an injective function of δ' , since $b''_{\max} = \alpha b'_{\max}$, $\alpha \in \mathbb{N}_+ \setminus \{1\}$.

To show that $A''_{c\delta''} < |C|$, we need to show that $A'_{c\delta'}b'_{\max} + \delta' < |C|b''_{\max}$. Observe that

$$\begin{aligned} A'_{c\delta'}b'_{\max} + \delta' &< |C|b'_{\max} + \delta' \\ &< |C|b'_{\max} + b'_{\max} = (|C| + 1)b'_{\max} \\ &\leq |C|\alpha b'_{\max} \\ &< |C|b''_{\max}. \end{aligned}$$

Finally, we have to show that $A''_{c\delta''} \neq A''_{c'\delta''}$ for $c \neq c'$. To see that, observe that if two elements from the matrix A' are transformed onto the same column in the matrix A'' , they must have been in the same column already in A' , since, as already mentioned earlier, $\delta'' = (A'_{c\delta'}b'_{\max} + \delta') \bmod b''_{\max}$ is an injective function of δ' . Consequently, for their preimages must hold $A'_{c\delta'} \neq A'_{c'\delta'}$. But then we obtain

$$\begin{aligned} &A'_{c\delta'} \neq A'_{c'\delta'} \\ \Leftrightarrow &\frac{(A'_{c\delta'}b'_{\max} + \delta') - \delta''}{b''_{\max}} \neq \frac{(A'_{c'\delta'}b'_{\max} + \delta') - \delta''}{b''_{\max}} \\ &A''_{c\delta''} \neq A''_{c'\delta''}. \end{aligned}$$

Finally, the remaining empty rows in each column of A'' can be filled arbitrarily, but such that for the added elements the two properties of a matrix representation still hold: $A''_{c\delta''} \neq A''_{c'\delta''}$ for $c \neq c'$, and $A''_{c\delta''} \in \{0, \dots, |C| - 1\}$. \square

5 Performance Analysis for GREEDY Algorithms

In this section, we formulate and prove two results on the optimality of the GREEDY algorithms. The first claim that we prove is that GREEDY algorithms generate listening schedules which are WDT-optimal (and consequently also complete) for BP sets from \mathbb{F}_2 . This result is stated in the following proposition.

Proposition 6. *For a BP set $B \in \mathbb{F}_2$, schedules generated by GREEDY are WDT-optimal.*

Proof. The idea for the proof is as follows. We first show that scanning a channel $c \in C$ in a time slot $t \in \{0, \dots, b_{\max}|C| - 1\}$ either results in the discovery of at least the configuration $(c, b_{\max}, t \bmod b_{\max})$, or it results in the discovery of 0 configurations. We then show that this implies that a GREEDY algorithm discovers all configurations $\{(c, b_{\max}, \delta) \mid c \in C, \delta \in \{0, \dots, b_{\max} - 1\}\}$ during the first $b_{\max}|C|$ time slots. Finally, we show that in \mathbb{F}_2 , this implies that all configurations are discovered.

Assume that scanning a channel $c \in C$ in a time slot $t \in \{0, \dots, b_{\max}|C| - 1\}$ does not result in the discovery of the configuration $(c, b_{\max}, t \bmod b_{\max})$. This implies that channel c has already been scanned during a time slot $t' \in \{ib_{\max} + t \bmod b_{\max} \mid i \in \{0, \dots, \lfloor t/b_{\max} \rfloor - 1\}\}$, so that the configuration $(c, b_{\max}, t \bmod b_{\max})$ has been discovered prior to the time slot t . However, this would imply that all other configurations $\{(c, b, t \bmod b) \mid b \in B \setminus \{b_{\max}\}\}$ have been discovered prior to the time slot t as well. This is because they also send their beacons during the time slot t' , since b_{\max} is an integer multiple of any $b \in B$. Consequently, if scanning a channel $c \in C$ in a time slot $t \in \{0, \dots, b_{\max}|C| - 1\}$ does not result in the discovery of the configuration $(c, b_{\max}, t \bmod b_{\max})$, it does not result in any discoveries.

At the same time, for each time slot $t \in \{0, \dots, b_{\max}|C| - 1\}$ there is at least one channel $c \in C$ that results in the discovery of a configuration $(c, b_{\max}, t \bmod b_{\max})$. If this would not be the case, all $|C|$ configurations $\{(c, b_{\max}, t \bmod b_{\max}) \mid c \in C\}$ would have had to be discovered prior to time slot t , that is, during one of the time slots $\{ib_{\max} + t \bmod b_{\max} \mid i \in \{0, \dots, \lfloor t/b_{\max} \rfloor - 1\}\}$. The number of such time slots, however, is strictly less than $|C|$, leading to a contradiction.

Thus, since a GREEDY algorithm maximizes the expected number of discoveries in each time slot, it would not select a channel that results in 0 discoveries if it can select a channel that results in at least one discoverable configuration. Consequently, in each time slot t it discovers one configuration from the set $\{(c, b_{\max}, t \bmod b_{\max}) \mid c \in C\}$. Since this set contains $b_{\max}|C|$ configurations, it implies that it takes $b_{\max}|C|$ time slots in order to discover all of them.

Note, however, that in \mathbb{F}_2 , a configuration (c, b, δ) , $c \in C$, $b \in B$, $\delta \in \{0, \dots, b - 1\}$, is discovered if the configuration (c, b_{\max}, δ) is discovered, since in each time slot during which the configuration (c, b_{\max}, δ) sends its beacons, also the configuration (c, b, δ) sends its beacons. \square

We remark that the assumption that the considered BP set is in \mathbb{F}_2 is crucial to prove the WDT-optimality, as shown by a counterexample in Example 2.

For the family of BP sets $\mathbb{F}_3 \subset \mathbb{F}_2$ we show the much stronger result that schedules generated by GREEDY are not only complete and WDT-optimal, but that they are recursive, and thus also MDT-optimal and NDoT-optimal. This result is formulated in the following proposition.

Proposition 7. *For BP sets from \mathbb{F}_3 , a schedule is GREEDY if and only if it is recursive. Consequently, for BP sets $B \in \mathbb{F}_3$ GREEDY schedules are complete, WDT-optimal, MDT-optimal, and NDoT-optimal.*

Proof. The claim that a recursive schedule is GREEDY for BP sets from \mathbb{F}_3 follows from Proposition 4, demonstrating that recursive schedules maximize the number of configurations detected until each time slot.

It remains to show that a GREEDY schedule is recursive. More precisely, we need to show that maximizing the expected number of discoveries in each time slot leads to discovering all configurations (c, b, δ) , $\delta \in \{0, \dots, b-1\}$, during the time slots $t \in \{0, \dots, b|C| - 1\}$.

Assume that this is not the case and that a configuration (c, b, δ) , $c \in C$, $b \in B$, $\delta \in \{0, \dots, b-1\}$, is discovered in a time slot $T_{(c,b,\delta)} \geq b|C|$. This implies that channel c was not scanned during any of the time slots $\{ib + \delta \mid i \in \{0, \dots, b-1\}\}$. Since there are $|C|$ such slots, this implies that there exists a different channel $c' \in C$ which was scanned during at least two of these time slots: $t', t'' \in \{ib + \delta \mid i \in \{0, \dots, b-1\}\}$. W.l.o.g., assume $t' < t''$.

Consider the two subsets $B', B'' \subseteq B$, with $B' = \{b' \in B \mid b' \leq b\}$, and $B'' = \{b'' \in B \mid b'' > b\}$. Observe that by scanning channel c during the time slot t'' at least the configurations $\{(c, b'', t'' \bmod b'') \mid b'' \in B''\}$ and the configuration $(c, b, t'' \bmod b)$ are detected. However, by scanning channel c' during the time slot t'' , at most the configurations $\{(c', b'', t'' \bmod b'') \mid b'' \in B''\}$ are detected. The reason that none of the configurations with a BP $b \in B'$ are detected during the time slot t'' is that they are all detected during the time slot t' , since b is an integer multiple of each $b' \in B'$.

Consequently, scanning channel c during time slot t'' would have resulted in larger expected number of discoveries than scanning channel c' . This leads to a contradiction since a GREEDY algorithm maximizes the expected number of discoveries in each time slot. \square

We remark that the assumption that the considered BP set is in \mathbb{F}_3 is necessary to prove the MDT-optimality of schedules generated by GREEDY, as shown by a counterexample in Example 2.

6 Performance Analysis for CHAN TRAIN

Proposition 8. *For a BP set $B \in \mathbb{F}_3$, CHAN TRAIN is GREEDY.*

Proof. Note that CHAN TRAIN does not perform a channel selection in each time slot but that it may skip several time slots if during those time slots the expected number of discoveries on the selected channel does not decrease. This may lead to a schedule that does not maximize the expected number of discoveries in each time slot, since it is possible that scanning a different channel during one of the skipped time slots would result in a higher value. This, however, would imply that the expected number of discoveries on that other channel has increased since the time slot when the current channel was selected. In the following, however, we show that for BP sets from \mathbb{F}_3 , the expected number of discoveries is monotonically decreasing over time.

W.l.o.g. we assume $B = \{b_0, \dots, b_{n-1}\}$, $n \in \mathbb{N}_+$, $b_i > b_j$ for $i > j$.

Observe that since over \mathbb{F}_3 schedules obtained by GREEDY are recursive, they discover all configurations with a BP $b \in B$ during the first $b|C|$ time slots. Note that since there are $b|C|$ such configurations and it is not possible to discover two of them during one time slot (they either have different channels or different offsets), exactly one such configuration must be detected during each of the first $b|C|$ time slots. Consequently, the number of configurations discovered during each time slot $t \in \{b_i|C|, \dots, b_{i+1}|C| - 1\}$ is $n - (i + 1)$. This means that the number of discoverable configurations is monotonically decreasing over time. Furthermore, since we assume that a neighbor selects

a certain channel c and beacon offset δ randomly according to a uniform distribution, it follows that the expected number of discoveries is also monotonically decreasing over time. \square

Note that for a BP set $B \notin \mathbb{F}_3$ Proposition 8 no longer holds, as illustrated by Example 3. However, CHAN TRAIN is still WDT-optimal on \mathbb{F}_2 .

Proposition 9. *For a BP set $B \in \mathbb{F}_2$, CHAN TRAIN is WDT-optimal.*

Proof. The proof proceeds along the lines of the proof for Proposition 6 using the observation that, analogously to GREEDY algorithms, CHAN TRAIN would never produce an idle slot as long as at least one discovery can be made. \square

7 Performance Analysis for OPT_{B2}

Proposition 10. *For arbitrary BP sets with $|B| = 2$, OPT_{B2} generates recursive listening schedules. Thus, they are complete, WDT-optimal, MDT-optimal, and NDoT-optimal.*

Proof. The claim follows directly from the definitions of OPT_{B2} and recursive schedules. \square

8 The WDT of MDT-Optimal Schedules

In this section, we will show that for an arbitrary set of BP's, a schedule minimizing the MDT uses at most $\text{LCM}(B)|C|$ time slots, and that, consequently, for BP sets in \mathbb{F}_2 , any MDT-optimal schedule is also WDT-optimal.

The following proposition establishes an upper bound on the number of time slots required to minimize an arbitrary strictly increasing function of discovery times. The idea for the proof is to show that any schedule that results in a configuration being detected after the time slot $\text{LCM}(B)|C|$ can be modified such that the configuration is detected before time slot $\text{LCM}(B)|C|$, without increasing the discovery times of other configurations.

Proposition 11. *For an arbitrary set of BP's $B \in \mathbb{F}_1$, a set of channels C , and a function $f : \mathbb{N}^{|K_{BC}|} \rightarrow \mathbb{R}$, which is strictly increasing in each argument, complete schedules \mathcal{L}^* that minimize $f((T_{\kappa}(\mathcal{L}))_{\kappa \in K_{BC}})$ have a WDT $T_{\mathcal{L}^*} \leq \text{LCM}(B)|C|$.*

Proof. Assume schedule \mathcal{L} minimizes f and $T_{\mathcal{L}} > \text{LCM}(B)|C|$. Consequently, there is at least one configuration $\kappa = (c, b, \delta)$ with discovery time $T_{\kappa}(\mathcal{L}) = T_{\mathcal{L}} > \text{LCM}(B)|C|$. Consider time slots $\tilde{\mathcal{T}}_{\kappa} = \{\delta + i\text{LCM}(B)\}_{i \in \{0, \dots, |C|-1\}}$. Observe that $(\{c\} \times \tilde{\mathcal{T}}_{\kappa}) \cap \mathcal{L} = \emptyset$ since otherwise κ would have been detected during one of the time slots in $\tilde{\mathcal{T}}_{\kappa}$. Consequently, there either exists an idle time slot $\tilde{t} \in \tilde{\mathcal{T}}_{\kappa}$, or, since $|\tilde{\mathcal{T}}_{\kappa}| = |C|$, there exist time slots $t', t'' \in \tilde{\mathcal{T}}_{\kappa}$, $t' \neq t''$, and a channel $c' \neq c$ such that $(c', t'), (c', t'') \in \tilde{\mathcal{T}}_{\kappa}$.

In the first case, we construct a new schedule $\mathcal{L}' = \mathcal{L} \setminus \{(c, T_{\mathcal{L}})\} \cup \{(c, \tilde{t})\}$, such that κ is detected during \tilde{t} and none of the discovery times of other neighbor configurations are increased.

In the second case, we construct a new schedule $\mathcal{L}' = \mathcal{L} \setminus \{(c, T_{\mathcal{L}}), (c', t'')\} \cup \{(c, t'')\}$. With the new schedule, configuration κ is detected during time slot t'' . In order to show that the discovery times of other neighbors do not increase, consider the vector function $\delta(t) = (t \bmod b, b \in B)$, providing for each time slot t a vector of offsets that can be detected in t . Since periodicity of $\delta(t)$ is $\text{LCM}(B)$,

we conclude that $\delta(t') = \delta(t'')$, and thus no discoveries are performed during time slot t'' with the schedule \mathcal{L} . Consequently, none of the discovery times are increased with the new schedule.

Repeating the above procedure for each κ with discovery time $T_\kappa(\mathcal{L}) > LCM(B)|C|$ results in a schedule \mathcal{L}^* with WDT $T_{\mathcal{L}^*} \leq LCM(B)|C|$ with $f((T_\kappa(\mathcal{L}^*))_{\kappa \in K_{BC}}) < f((T_\kappa(\mathcal{L}))_{\kappa \in K_{BC}})$, proving the claim. \square

The following Corollary presents a notable consequence from Proposition 11 for BP sets from \mathbb{F}_2 .

Corollary 12. *For a BP set $B \in \mathbb{F}_2$, a set of channels C , and a function $f : \mathbb{N}^{|K_{BC}|} \rightarrow \mathbb{R}$, which is strictly increasing in each argument, complete schedules \mathcal{L}^* that minimize $f((T_\kappa(\mathcal{L}))_{\kappa \in K_{BC}})$ are WDT-optimal.*

Proof. The claim follows directly from the definition of WDT-optimality, Proposition 11, and the defining property of $B \in \mathbb{F}_2$ that $LCM(B) = \max(B)$. \square

Please observe that the fact that MDT satisfies the conditions on the function f in Proposition 11 implies that the upper bounds established in Proposition 11 and Corollary 12 also apply to schedules minimizing MDT.

Corollary 13. *For an arbitrary set of BP's $B \in \mathbb{F}_1$, and a set of channels C , MDT-optimal listening schedules \mathcal{L}^* have a WDT $T_{\mathcal{L}^*} \leq LCM(B)|C|$.*

Proof. Observe that MDT is strictly increasing in each configuration detection time T_κ . Applying Proposition 11 we obtain the claim. \square

Corollary 14. *For a set of BP's $B \in \mathbb{F}_2$, and a set of channels C , MDT-optimal listening schedules are also WDT-optimal.*

Proof. Observe that MDT is strictly increasing in each configuration detection time T_κ . Applying Corollary 12 we obtain the claim. \square

9 NDoT-Optimality

Proposition 15. *A NDoT-optimal schedule \mathcal{L} is also WDT-optimal.*

Proof. Since a NDoT-optimal schedule pointwise maximizes the CDF of discovery times, the value of the CDF for the normalized discovery time 1 is also maximized. Since there always exist schedules that discover all neighbors at the normalized discovery time 1, and therefore whose CDF's at normalized time 1 have the value 1, we conclude that the CDF of a NDoT-optimal schedule also must have the value 1 at the normalized time 1. \square

Proposition 16. *If for a set of BP's $B \in \mathbb{F}_1$ and a set of channels C a NDoT-optimal schedule \mathcal{L}^* exists, it follows that any MDT-optimal schedule \mathcal{L} is also NDoT-optimal.*

Proof. The proof is constructed by contradiction. Assume that schedule \mathcal{L}^* is NDoT-optimal. It implies that \mathcal{L}^* is also MDT-optimal. Furthermore, assume that another schedule \mathcal{L} is MDT-optimal but not NDoT-optimal. The CDF's of normalized discovery times of \mathcal{L} and \mathcal{L}^* shall be denoted by f^* and f , respectively. Since both schedules are MDT-optimal we obtain:

$$\int_0^{T^*} x f^{*'}(x) dx = \int_0^T x f'(x) dx,$$

where $T^*, T \geq 1$ denote the respective WDT's.

Applying partial integration and using (i) the fact that $T^* = 1$ and $f^*(1) = 1$ due to the WDT-optimality of \mathcal{L}^* , and (ii) the fact that $f(0) = f^*(0) = 0$, we obtain

$$\begin{aligned} & x f^*(x) \Big|_0^1 - \int_0^1 f^*(x) dx - x f(x) \Big|_0^b + \int_0^b f(x) dx = 0 \\ & 1 - \int_0^1 f^*(x) dx - b + \int_0^b f(x) dx = 0 \\ & \Rightarrow 1 - b - \int_0^1 f^*(x) dx + \int_0^1 f(x) dx + \int_1^b f(x) dx = 0 \\ & \Rightarrow (b-1) + \left[\int_0^1 f^*(x) dx - \int_0^1 f(x) dx \right] - \int_1^b f(x) dx = 0 \end{aligned}$$

On the other hand, since \mathcal{L} is not NDoT-optimal: $\int_0^1 f^*(x) dx - \int_0^1 f(x) dx > 0$. Furthermore, observe that $\int_1^b f(x) dx \leq b-1$. We obtain a contradiction proving the original claim. \square

10 Computational Complexity and Memory Requirements

Note that GREEDY algorithms have a compelling property of low complexity. A possible implementation of pseudo code in Algorithm 1 presented in [5] proceeds as follows. For each configuration it stores a binary variable indicating if it has been covered or not, resulting in $|C| |B| \sum_{b \in B} b$ bits of required memory space. At each time slot t , it iterates over all channels $c \in C$, computing for each channel, which of configurations $\{(c, b, t \bmod b) \mid b \in B\}$ are not yet considered. Finally, it selects a channel, for which the sum of probabilities of these configurations to be selected by a neighbor is the highest. Consequently, computational complexity at each time slot is $\mathcal{O}(|C| |B|)$. It iterates over time slots until all configurations are covered.

The overall computational complexity depends on the number of time slots required to consider all configurations. Since GREEDY algorithms are WDT-optimal for $B \in \mathbb{F}_2$, the number of required time slots is $\max(B) |C|$, resulting in a total complexity of $\mathcal{O}(|C|^2 |B| \max(B))$. For BP sets from \mathbb{F}_1 , GREEDY algorithms require at most $LCM(B) |C|$. A proof of this claim is similar to the proof of Proposition 6 and is omitted for brevity.) Thus, an upper bound for the computational complexity over \mathbb{F}_1 is $\mathcal{O}(|C|^2 |B| LCM(B))$.

Please note that if the assumption made in our system model in Section III in [5] that the Greatest Common Divisor (GCD) of the considered BP sets is 1 does not hold, the complexity can be further reduced by replacing B by $B' = \left\{ \frac{b}{\text{GCD}(B)} \right\}_{b \in B}$. This preprocessing step allows to reduce computational complexity over \mathbb{F}_2 to $\mathcal{O}(|C|^2 |B| \frac{\max(B)}{\text{GCD}(B)})$, over \mathbb{F}_1 the upper bound becomes $\mathcal{O}(|C|^2 |B| \frac{LCM(B)}{\text{GCD}(B)})$.

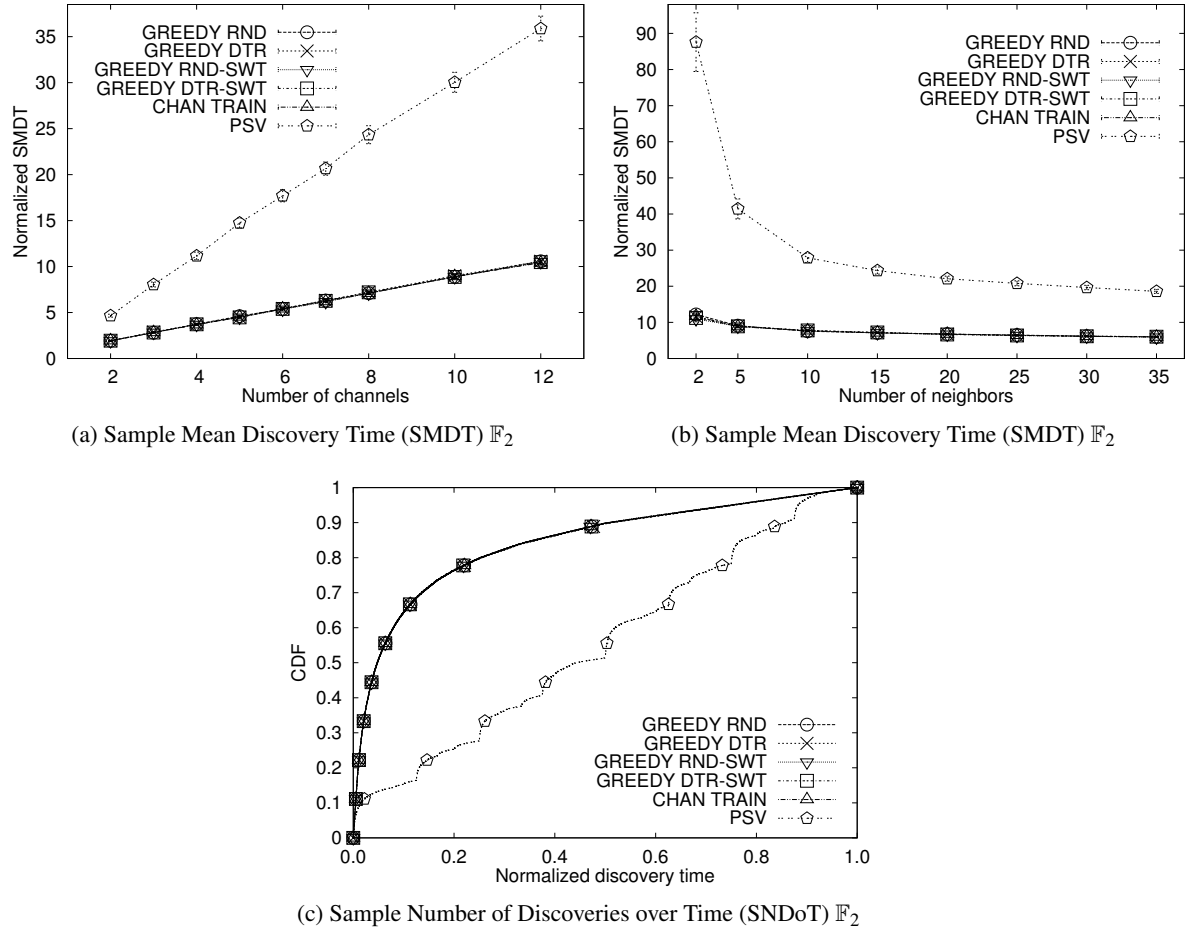


Figure 6: Evaluation results obtained by simulation for the family of BP sets \mathbb{F}_2 (see Section 11.2).

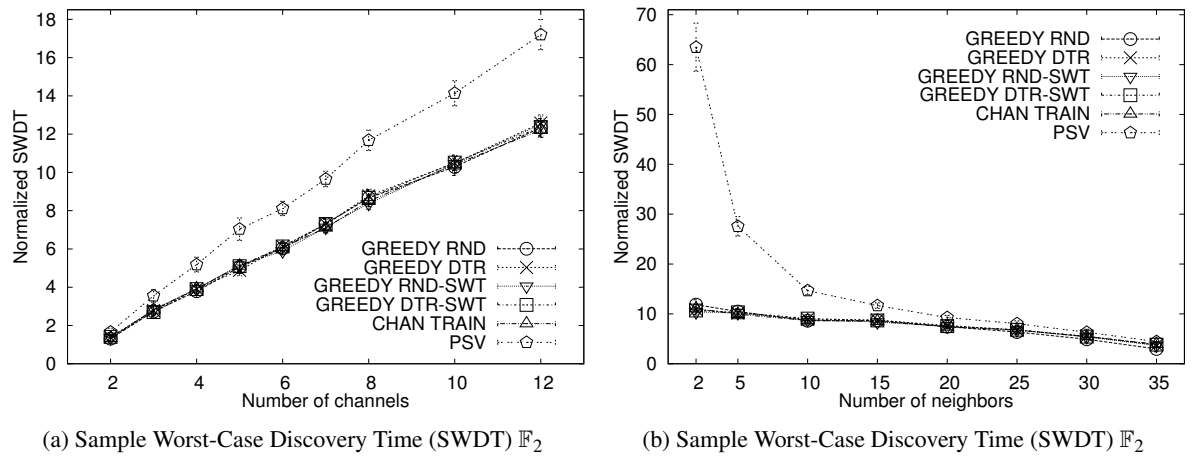


Figure 7: Evaluation results obtained by simulation for the family of BP sets \mathbb{F}_2 (see Section 11.2).

11 Additional simulation results

In the following we present supplementary simulation results evaluating three performance metrics that are only defined for specific instances of the studied problem, and are therefore named Sample Mean Discovery Time (SMDT), Sample Number of Discoveries over Time (SNDOT) and Sample Worst-Case Discovery Time (SWDT). These metrics are more precisely specified in the following. The setting is described in Section IX-A in [5].

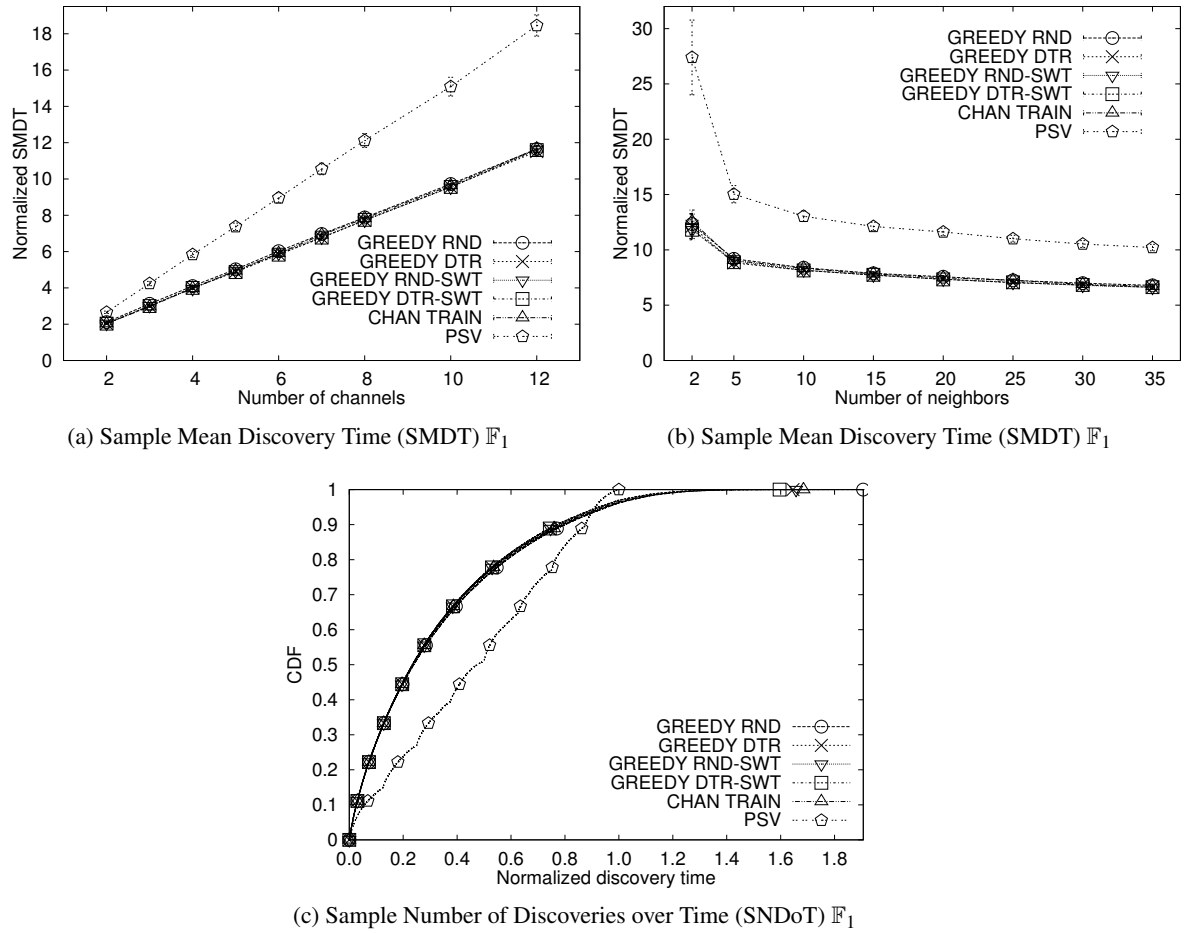
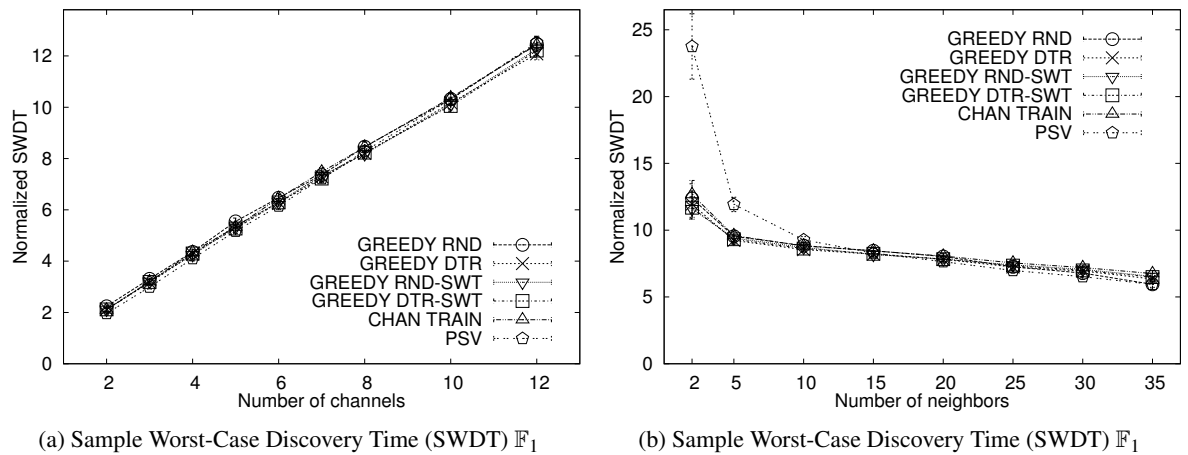
11.1 Performance Metrics

In the following we evaluate three performance metrics: SNDOT, SMDT, and SWDT. The SMDT is defined as $\frac{1}{|N|} \sum_{v \in N} T_v(\mathcal{L})$. It relates to the MDT in the way a sample mean relates to the 'true' mean. Since a sampled mean is an unbiased estimator, the expected value of the SMDT is the MDT. The SWDT is defined as $\max_{v \in N} T_v(\mathcal{L})$. It is the analogon of the WDT in a specific network environment, where only a subset of possible configurations is present. However, while the SMDT can be interpreted as an estimation of MDT, SWDT has no such relationship with the WDT. In fact, the WDT is a more important property of a listening schedule, since it determines the required execution time in scenarios, where a complete discovery is desired. The SWDT, on the other hand, cannot be used as an indicator for the discoverer to stop the discovery process, since the discoverer cannot know if all neighbors have been discovered or not, before it has probed for all potential configurations, which requires an execution time equal to the WDT. Still, when the number of neighbors increases, SWDT converges to the WDT due to the law of large numbers. Finally, SNDOT is the number of discovered neighbors divided by the total number of neighbors as a function of the execution time of the schedule. Also for the SNDOT the law of large numbers implies convergence to the CDF of discovery times for large numbers of neighbors.

The discovery times used in SNDOT are normalized as described in Section IV in [5]. SMDT and SWDT are normalized w.r.t. their corresponding optimum values obtained by solving an ILP similar to MDTOPT but with a different objective function. For each scenario with a given set of neighbors and their configurations two ILP's are solved that generate listening schedule minimizing the SMDT and SWDT among those neighbors. We use the following optimization variables.

$$y_{iv} = \begin{cases} 1, & \text{if neighbor } v \in N \text{ is detected during time slot} \\ ib_v + \delta_v & \\ 0, & \text{otherwise} \end{cases}$$
$$h_{ct} = \begin{cases} 1, & \text{if channel } c \text{ is scanned during time slot } t \\ 0, & \text{otherwise} \end{cases}$$

We use the following constraints.

Figure 8: Evaluation results obtained by simulation for the family of BP sets \mathbb{F}_1 (see Section 11.2).Figure 9: Evaluation results obtained by simulation for the family of BP sets \mathbb{F}_1 (see Section 11.2).

$$\sum_{i=0}^{\lfloor \frac{t_{max}-\delta_v}{b_v} \rfloor} y_{iv} = 1 \quad \text{for all } v \in N \quad (\tilde{C}1)$$

$$y_{iv} \leq h_{c_v, \delta_v + ib_v} \quad \text{for all } v \in N, \\ i \in \left\{ 0, \dots, \lfloor \frac{t_{max}-\delta_v}{b_v} \rfloor \right\} \quad (\tilde{C}2)$$

$$\sum_{c \in C} h_{ct} \leq 1 \quad \text{for all } t \in \{0, \dots, t_{max}\} \quad (\tilde{C}3)$$

Constraint $\tilde{C}1$ ensures that for each neighbor $v \in N$ a scan is scheduled in exactly one of the beacon time slots \mathcal{T}_v . The last two constraints $\tilde{C}2$ and $\tilde{C}3$ make sure that during each time slot not more than one channel is scanned. In comparison to the numerical evaluation we allow for much larger BP's from $B \in \mathbb{F}_1$. Therefore, to keep the ILP's computationally feasible, the maximum time slot used in the optimization problems is limited to $t_{max} = \min(1000 \max(B)|C| - 1, \text{LCM}(B)|C| - 1)$.

We use the following objective function to find the optimum SMDT for a given network environment.

$$\min \frac{1}{|N|} \sum_{v \in N} \sum_{i=0}^{\lfloor \frac{t_{max}-\delta_v}{b_v} \rfloor} y_{iv} (ib_v + \delta_v).$$

We use the following objective function, together with an additional constraint, to find the optimum SWDT for a given network environment.

$$\begin{aligned} \min \quad & z \\ \text{s.t.} \quad & y_{iv} (ib_v + \delta_v) \leq z \\ & \text{for all } v \in N, i \in \left\{ 0, \dots, \lfloor \frac{t_{max}-\delta_v}{b_v} \rfloor \right\}. \end{aligned}$$

Note that the results for the SMDT and SWDT presented in the following cannot be directly compared with the numerical results of the MDT and WDT presented in Section VIII-B in [5] due to different methods of normalization.

11.2 Results

In the following, we present simulations results for BP families \mathbb{F}_1 and \mathbb{F}_2 accompanied by confidence intervals for a confidence level of 95%.

Results for \mathbb{F}_2

Figures 6a and 6b depict the results for the normalized SMDT, as a function of the number of channels and of the number of neighbors, respectively. For two channels the PSV strategy results in a SMDT which is about twice as high as the results of the GREEDY approaches and CHAN TRAIN. However, the gap is increasing with larger number of channels. For twelve channels the SMDT of PSV is more

than 350% higher compared to the results of GREEDY / CHAN TRAIN. For all analyzed number of neighbors PSV results again in a MDT which is at least 350% higher than the results achieved by GREEDY and CHAN TRAIN. In addition, all strategies diverge from the optimum with an increasing number of channels, while they approach the optimum with an increased number of neighbors.

Figure 6c shows the results for SNDOT. GREEDY and CHAN TRAIN discover up to 50% more neighbors in average at certain points in time without having an impact on the discovery time of the last percentage of discovered neighbors.

Figures 7a and 7b display the normalized SWDTs. With two channels all strategies result in similar values. However, for larger number of channels, and even more so for low numbers of neighbors, the PSV strategy results in a significantly larger SWDT. Again, we remark that the SWDT is of a lower importance than the WDT studied in Section VIII in [5] (see also Section IV).

Results for \mathbb{F}_1

Figures 8 and 9 show the results for the family of BP sets \mathbb{F}_1 . The findings are qualitatively mostly similar to those for BP sets from \mathbb{F}_2 .

The normalized SMDT as a function of the number of channels is depicted in Figures 8a and 8b. Similar to the results for \mathbb{F}_2 our approaches significantly outperform PSV. However, the gap between them has decreased.

Figure 8c shows the results of the SNDOT. The majority of neighbors is discovered faster by the GREEDY strategies and CHAN TRAIN. However, about 10% of neighbors are discovered later by GREEDY and CHAN TRAIN as compared to PSV.

The normalized SWDT is shown in Figures 9a and 9b. In contrast to the results for \mathbb{F}_2 the PSV strategy achieves similar results as our strategies for all analyzed number of channels.

11.3 Summary

The simulation results reveal that even under realistic conditions the proposed algorithms significantly outperform the PSV strategy regarding the SMDT for the family of BP sets \mathbb{F}_2 as well as \mathbb{F}_1 and SWDT for \mathbb{F}_2 for different number of channels and number of neighbors. The results for SNDOT are similar to the numerical results presented in Section VIII-B in [5]. For the family of BP sets \mathbb{F}_2 our strategies discover up to 50% more neighbors compared to PSV at certain points in time without penalizing the discovery of the last neighbors. However, for the family of BP sets \mathbb{F}_1 the gap decreases and the discovery of the last 10% of neighbors requires additional time by our approaches in comparison with PSV.

Bibliography

- [1] “Gartner Press Release,” <http://www.gartner.com/newsroom/id/2905717>.
- [2] D. Evans, “The Internet of Things - How the Next Evolution of the Internet Is Changing Everything,” Cisco Internet Business Solutions Group (IBSG), Tech. Rep., April 2011.
- [3] “IEEE Standard 802.11-2012,” 2012.
- [4] “IEEE Standard 802.15.4-2006,” 2006.
- [5] N. Karowski, K. Miller, and A. Wolisz, “Greedy multi-channel neighbor discovery,” vol. abs/1807.05220, 2018. [Online]. Available: <http://arxiv.org/abs/1807.05220>
- [6] D. B. West, *Introduction to Graph Theory*, 2nd ed. Pearson, 2001.