

ON OPTIMAL NEIGHBOR DISCOVERY

PHILIPP H. KINDT AND SAMARJIT CHAKRABORTY

ABSTRACT. Mobile devices apply neighbor discovery (ND) protocols to wirelessly initiate a first contact within the shortest possible amount of time and with minimal energy consumption. For this purpose, over the last decade, a vast number of ND protocols have been proposed, which have progressively reduced the relation between the time within which discovery is guaranteed and the energy consumption. In spite of the simplicity of the problem statement, even after more than 10 years of research on this specific topic, new solutions are still proposed even today. Despite the large number of known ND protocols, given an energy budget, what is the best achievable latency still remains unclear. This paper addresses this question and for the first time presents safe and tight, duty-cycle-dependent bounds on the worst-case discovery latency that no ND protocol can beat. Surprisingly, several existing protocols are indeed optimal, which has not been known until now. We conclude that there is no further potential to improve the relation between latency and duty-cycle, but future ND protocols can improve their robustness against beacon collisions.

1. INTRODUCTION

Wireless networks that operate without any fixed infrastructure are rapidly growing in importance. Since all devices in such mobile ad-hoc networks (MANETs) run on batteries or rely on intermittently available energy-harvesting sources, the energy spent for communication needs to be as low as possible. Typically, MANET radios are duty-cycled and wake up only for short durations of time for carrying out the necessary communication and then go back to a sleep mode. While such duty-cycled communication schemes are easy to realize when the clocks of all devices are synchronized and their wakeup schedules are known by all participants of the network, asynchronous communication (i.e., communication without synchronized clocks) remains a challenging problem. One of the most important asynchronous procedures is establishing the first contact between different wireless devices, which is referred to as *neighbor discovery (ND)*.

Neighbor Discovery: ND is used by a device for detecting other devices in range. This could be for clock synchronization and establishing a connection, after which more data can be exchanged in a synchronous fashion. Efficient ND is characterized by achieving the shortest possible discovery latency for a given energy budget. Towards this, a large number of ND protocols have been proposed till date, see [1, 2, 3, 4, 5, 6, 7, 8, 9, 10, 11, 12, 13, 14, 15, 16, 17, 18, 19, 20, 21, 22, 23, 24, 25, 26, 27, 28, 29, 30, 31, 32, 33, 34, 35, 36, 37, 38, 39, 40, 41, 42, 43, 44, 45, 46, 47, 48]. Among these, [5, 6, 7, 8, 9, 10, 11, 12, 13, 14, 15, 16, 17, 18] for example, concern *deterministic* discovery. Here, given the protocol parameters, an upper bound on the discovery latency can be determined. The problem of *pairwise discovery* between two devices is of fundamental importance, since in many scenarios, devices join the network gradually and only a master device and the newly joining one carry out the discovery procedure simultaneously. Moreover, the process of discovering multiple devices always relies on pairwise ND.

Over the years, successive ND protocols have improved their discovery latencies for given energy budgets. For example, the *Griassdi* [17] protocol proposed in 2017 claims to achieve by 87% lower worst-case latencies than

The authors are with the Chair of Real-Time Computer Systems (RCS), Department of Electrical and Computer Engineering, Technical University of Munich (TUM), Germany. Email: philipp.kindt/samarjit[at]tum.de.

Searchlight-Striped [9] that was proposed in 2012. However, despite the significant attention the ND problem has received over the past 15+ years, the fundamental question of what is the theoretically lowest possible discovery latency that any ND protocol could guarantee for a given energy budget still remains unanswered.

Performance of ND Protocols: In the absence of such a bound, the performance evaluations of different ND protocols have often been very subjective. The results of such evaluations relied on the choice of protocols, their parametrizations and the assumed setups. Hence, while a certain protocol might outperform others in such a comparison, it might perform differently if the parametrization or setup is changed. In addition, most known protocols, e.g., [7, 9, 14], subdivide time into multiple slots and are hence referred to as *slotted*. The device sleeps in most slots, whereas some slots are active and used for communication. In each active slot, a device sends a beacon at the beginning and/or end of the slot and listens for incoming beacons in the meantime. Discovery occurs once two active slots overlap in time. Here, performance is quantified in terms of the worst-case number of slots until discovery is guaranteed. Though a certain protocol could perform better than another in terms of the number of slots, such comparisons are heavily dependent on the supported range of slot lengths. As a result, such comparisons in terms of slots and not directly in terms of time are often not meaningful. Moreover, despite slotted protocols having been studied thoroughly in the literature, many protocols that are frequently used in practice, e.g., Bluetooth Low Energy (BLE), do not rely on a slotted paradigm. They schedule reception windows and beacon transmissions with periodic intervals and offer three degrees of freedom that can be configured freely (viz., the periods for reception and transmission, and the length of the reception window). The high practical relevance of such *periodic interval (PI)-based* protocols is underpinned by the 4.7 billion BLE-devices that were expected to be sold in 2018 [49]. It has recently been shown that the parametrizations for ND in BLE networks proposed by official specifications [50] lead to performances far from the optimum [18]. This has raised the interest to fully understand such *slotless* ND procedures. In particular, finding beneficial parametrizations for periodic interval-based protocols has been studied in the literature recently, e.g., in [18, 16, 17]. However, until today, it is neither clear whether the proposed parametrizations are actually optimal, nor how such protocols compare to the slotted ones in terms of performance. In summary, despite the large volume of available literature, it is not possible to meaningfully assess and classify the performance of ND protocols in a purely objective fashion.

This Paper: In this paper, we study the fundamental limits of pairwise, deterministic ND. In particular, we establish a relationship between the optimal discovery latency, channel utilization (and hence beacon collision rate) and duty-cycle. No pairwise ND protocol can achieve lower discovery latencies than the ones established in this paper. The resulting bounds not only give important insights into the design of ND protocols, but will serve as a baseline for more objective performance comparisons. Surprisingly, our analysis shows that some recently proposed protocols actually perform optimally and cover the entire discovery latency/channel utilization/duty-cycle Pareto front. The optimality results of such protocols were not known until now. In particular, the coverage of the Pareto front implies that there is no further potential for improvement. However, there is still potential to improve the robustness against beacon collisions, which might occur frequently when many devices carry out ND simultaneously.

Principle of ND: In general, a radio can either be in a sleep state, listen to the channel or transmit a beacon. Hence, the basic building blocks of a ND protocol are given by these three operations and any ND protocol can be represented as a unique pattern of them. For a higher power-budget, the number of beacons and/or the number or lengths of reception windows can be increased and a discovery procedure is successful once a beacon overlaps with a reception window on another device. Since the design space of all possible reception and transmission patterns allows for an infinite number of possible configurations, determining the optimal pattern and its performance through any form of exhaustive search or numerical method is not possible. Further, as outlined above, most work on ND has focused on slotted protocols and therefore studied only a small part of the design space. As a result, the problem of assessing the optimal performance of ND has so far remained unsolved.

ND Scenarios: For different scenarios, the ND problem appears in different forms, and we provide bounds on the discovery latency for many of them. First, it is obvious that if two devices E and F both have the same beacon and reception patterns, their discovery properties are *symmetric*. This implies that device E discovers device F with the same worst-case latency for a given duty-cycle as F discovering E . Several publications, e.g., [21, 7, 18], have studied this special case of *symmetric* duty-cycles, for which we present a bound on the discovery latency. If both devices run different patterns (for example, due to different duty-cycles), the discovery properties are *asymmetric*. For the asymmetric case, we provide a bound on the discovery latency when each device is aware of the other device’s configuration.

Another important question we answer in this paper is the partitioning of the duty-cycle, which corresponds to the energy-budget of a device. The duty-cycle of a device is the fraction of time it is active. On the other hand, channel utilization is the fraction of time a device occupies the channel, which is between zero and its duty-cycle. Beacon collision rates are solely determined by the channel utilizations of the devices in range. For the case when the channel-utilization (and hence collision rate) is unconstrained, we derive the ratio between transmission and reception times that minimizes the discovery latency.

In the case of many devices discovering each other, the channel utilization of each device has to be constrained for limiting the collision rate. In this paper, we therefore not only derive bounds for the discovery latency that any protocol can guarantee for a given duty-cycle, but also for the case where both duty-cycle and the maximum channel-utilization are provided. To the best of our knowledge, no such protocol-agnostic bounds on discovery latency for the ND problem has been derived until now. In particular, this paper makes the following contributions.

Technical Contributions: We present the following bounds on the discovery latency of deterministic ND protocols.

- (1) The lowest discovery latency *any* symmetric and asymmetric pairwise ND protocol can guarantee for a given duty-cycle and hence energy consumption. Recall that in symmetric ND, all devices operate using the same duty-cycle, whereas in asymmetric ND devices use different duty-cycles.
- (2) A discovery latency bound for the case where the channel utilization is additionally constrained.
- (3) Bounds for the following three cases where two devices E and F discover each other. (a) Only E needs to discover F , whereas F does not need to discover E . (b) Either E discovers F or F discovers E , but both discovering each other is not possible. (c) Both E and F mutually discover each other.

We further study the relation between our bounds and previously known ones [20, 21, 10, 11], which are all limited to slotted protocols. These bounds are given in terms of a worst-case number of slots until discovery is guaranteed, where the discovery latency also depends on the slot length. However, how small a slot length can be is difficult to answer, while it is known that slot lengths cannot be made arbitrarily small. Therefore, discovery latencies in terms of time have not been derived, which we address in this paper. Finally, while most previous work has focused on slotted protocols, we show that when channel utilization is unconstrained, only slotless protocols can perform optimally, whereas slotted ones cannot. This result is important because in many IoT scenarios devices join gradually and only a pair of devices participate in ND at any point in time. Here, channel utilization is therefore not of concern.

Organization of the paper: The rest of this paper is organized as follows. In Section 2, we present related work on discovery latency bounds of ND protocols. Next, in Section 3, we provide a formal description of a generic ND procedure. Based on this, in Section 4, we derive a list of properties that deterministic ND protocols need to guarantee. Recall that deterministic ND protocols are ones for which bounded discovery latencies can be guaranteed. We derive such latency bounds in Section 5. Finally, in Section 7, we relate the latency bounds of multiple existing ND protocols to the bounds obtained in this paper. Throughout this paper, we make a couple of simplifying assumptions. These assumptions are only for the ease of exposition and are relaxed in Section 6. A table of symbols and additional proofs are given in the appendix.

2. RELATED WORK

In this section, we describe existing efforts to determine bounds on the discovery latency that any ND protocol can achieve, and relate them to this paper.

Bounds for Slotted Protocols: As discussed above, the vast majority of ND protocols proposed in the literature follow a slotted paradigm, in which reception and transmission are temporally coupled into slots. A bound on the discovery latency of slotted protocols has been studied in [20, 21]. Here, it has been shown that for guaranteeing discovery within T slots, every device needs to have at least $k = \sqrt{T}$ active slots. Therefore, if e.g., $k = 2$ out of $T = 4$ slots are active, then discovery can be guaranteed within four slots with a duty-cycle of 50%, whereas if $k = 4$ and $T = 16$, discovery can be guaranteed within 16 slots with a duty-cycle of 25%. Determining the schedule of active slots that realizes this bound relies on *cyclic difference sets* [20]. Since only a very limited number of such difference sets are known, slotted protocols utilizing this bound can only be realized for a few duty-cycles that correspond to these known difference sets. Subsequently proposed protocols, such as Disco [7], Searchlight [9] and U-Connect [8] for the same discovery latency require more active slots than defined by this bound. But they are more flexible in terms of duty-cycles they can realize. Other recent work [10, 11] claims to have superseded this bound. By sending an additional beacon outside the slot boundaries in a schedule defined by difference sets, a tighter bound than described in [20, 21] can be reached.

Being on slotted protocols, the bounds in [20, 21, 10, 11] are all given in terms of a worst-case number of slots within which discovery is guaranteed. The corresponding bounds in terms of time depend on the slot length I . The minimum size of such a slot, among other factors, also depends on the hardware, and cannot be made arbitrarily small. Consequently, no bounds on the discovery latency in terms of time for slotted protocols have been known until now. This issue is addressed later in this paper.

Bounds for PI-based Protocols: Given a tuple of parameter values (T_a, T_s, d_s) , a method to compute the worst-case discovery latency for PI-based protocols was provided in [51]. However, since there are infinite numbers of possible parametrizations (T_a, T_s, d_s) , and because of the computation scheme provided in [51], which parametrization leads to the lowest discovery latency has so far remained unknown. Recently, [18] and [17] proposed parametrization schemes that can compute parameters (T_a, T_s, d_s) to realize any given duty-cycle. However, the optimality of such parametrizations in terms of discovery latency has not been established.

Generic Approaches: Unlike the work described above that was specific to slotted or PI-based protocols, protocol-agnostic bounds were presented in [52, 53]. In particular, they give an asymptotic latency bound in the form of $\Theta(d)$, where d is “the discretized uncertainty period of the clock shift between the two processors” [52]. Hence, this bound depends on the degree of asynchrony between the clocks of a sender and a receiver. First, the asymptotic nature of such a bound is very different from the concrete time bounds that have been pursued within the computer communications community, e.g., [20, 21, 10, 11], and the ones presented by us in this paper. Second, this community has also focused on bounds that depend on duty-cycle and hence energy budget, which are of direct practical relevance. For these reasons, the bounds from [52, 53] are not comparable to those that have been more commonly pursued, and also to those presented in this paper.

3. NEIGHBOR DISCOVERY PROTOCOLS

3.1. Definition. In this section, we formally define the ND procedure and its associated properties.

Definition 3.1 (Reception Window Sequence): Let the time windows during which a device listens to the channel be given by the tuples $c_1 = (t_1, d_1), c_2 = (t_2, d_2), c_3 = (t_3, d_3), \dots$, where each *reception window* c_i starts at time t_i and ends d_i time-units later (see Figure 1a)). A *reception window sequence* $C = c_1, c_2, \dots, c_n$ could be of finite or infinite length. In this paper, for simplicity of notation, we refer to such finite length sequences by C and infinite length sequences by C_∞ .

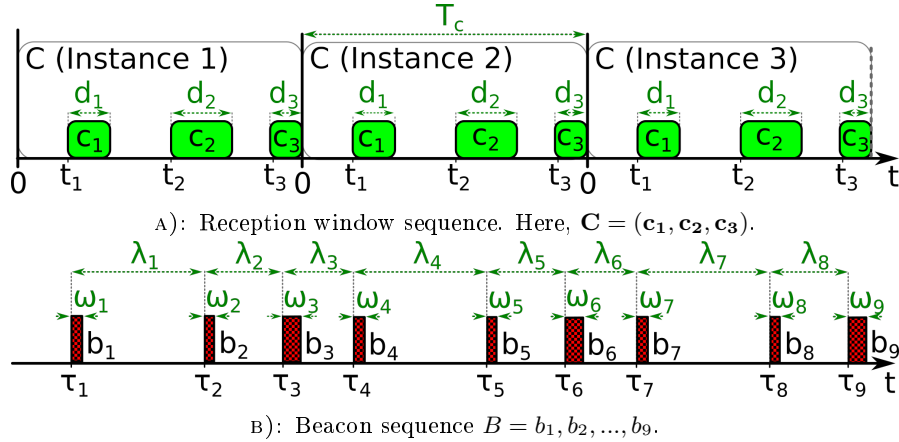


FIGURE 1. Reception window and beacon sequences.

For the simplicity of exposition, throughout this paper, we always assume that any C_∞ is an infinite concatenation of some finite length sequence C . For such C_∞ , we define $n_C = |C|$ (i.e., the number of windows contained in C). Further, we denote the time between the ends of two consecutive instances of C as the *reception period* T_C . It is worth mentioning that all our bounds remain valid also for sequences C_∞ that are not given by concatenating the same C , as we show in Appendix A. We assign a time-axis to every instance of C . For convenience, which will become clear later, the origin of time in a certain instance of C will start at the end of the last reception window of the previous instance, as depicted in Figure 1a). In this figure, C consists of three reception windows (i.e., c_1, c_2, c_3), and three concatenated instances of C are shown. For example, the origin of the time-axis for Instance 2 lies at the end of c_3 in Instance 1.

Definition 3.2 (Beacon Sequence): A sequence of beacons $B = b_1, b_2, \dots, b_m$ sent at the time-instances $\tau_1, \tau_2, \dots, \tau_m$, as depicted in Figure 1b), is called a *beacon sequence* of length m . The transmission durations of these beacons are given by $\omega_1, \omega_2, \dots, \omega_m$. A sequence of infinite length (i.e., $m \rightarrow \infty$) is denoted by B_∞ .

We denote infinite length beacon sequences B_∞ that are given by concatenations of a finite beacon sequence B as *repetitive beacon sequences*. In such repetitive sequences, $m_B = |B|$ and the time between the endings of two consecutive instances of B is given by T_B . Unlike for reception window sequences, we do not restrict ourselves to repetitive infinite beacon sequences. However, we will prove that all beacon sequences that optimize the relevant metrics of a ND procedure are repetitive when the corresponding reception window sequence is also repetitive.

We indicate an arbitrary shorter sequence B' to be a part of a longer sequence B by using the notation $B' \in B$. For example, in Figure 1b), $B' = b_2, b_3, b_4, b_5, b_6 \in B$. Further, the time between the beginnings of beacon b_i and beacon b_{i+1} is called the *beacon gap* λ_i . It is $\lambda_i = \tau_{i+1} - \tau_i$.

Definition 3.3 (ND Protocol): A tuple of an infinite beacon and reception window sequence (B_∞, C_∞) is called a *ND protocol*.

In this paper, unless explicitly stated, we assume that C_∞ and B_∞ stem from two different devices E and F . When it is necessary to explicitly specify the device that a sequence is scheduled on, we use the notation $C_{E,\infty}$ or $B_{F,\infty}$, where E and F refer to device E or F respectively. We also apply this notation to reception windows and beacons, e.g., $b_{E,1}$ refers to beacon 1 on device E and $c_{F,1}$ refers to reception window 1 on device F .

The most important properties of a ND protocol are its worst-case latency L , its duty-cycle η , and its channel utilization β , as defined next.

Definition 3.4 (Worst-Case Latency): Given two devices E and F , where E runs an infinite beacon sequence and F an infinite reception window sequence, the *worst-case latency* L is the earliest possible time after which




FIGURE 2. a) Beacons starting within the hatched area of window c_i can only fractionally coincide. b) Offset Φ_1 of the first beacon b_1 and Φ_2 of the second beacon b_2 in range.

an overlap of a beacon from E with a reception window of F is guaranteed, measured from the point in time both devices come into the range of reception.

Definition 3.5 (Duty-Cycle): The *transmission duty-cycle* β of a device is the fraction of time it spends for transmission, whereas the *reception duty-cycle* γ is the fraction of time spent for reception. In general, depending on the configuration of the radio (e.g., transmit power and receiver gain), transmission incurs a different power consumption than reception. Therefore, the total *duty-cycle* η is given as a weighted sum $\eta = \gamma + \alpha\beta$, where the weight α is the ratio of transmission and reception powers, i.e., $\alpha = P_{Tx}/P_{Rx}$. For a radio running a tuple of sequences (B_∞, C_∞) , it is:

$$(1) \quad \beta = \lim_{m \rightarrow \infty} \frac{\sum_{i=1}^{m-1} \omega_i}{\tau_m - \tau_1}, \quad \gamma = \lim_{n \rightarrow \infty} \frac{\sum_{i=1}^{n-1} d_i}{t_n - t_1}, \quad \eta = \alpha\beta + \gamma$$

The transmission duty-cycle β is the same as the *channel utilization*. The duty-cycle η directly corresponds to the power consumption of an ideal radio. Non-ideal radios are discussed in Section 6.3.

It follows from the above that the duty-cycle of a tuple of sequences B_∞, C_∞ , that are concatenations of finite length sequences B and C respectively, can be computed as follows.

$$(2) \quad \beta = \frac{\sum_{i=1}^{m_B} \omega_i}{T_B} = \frac{\sum_{i=1}^{m_B} \omega_i}{\sum_{i=1}^{m_B} \lambda_i}, \quad \gamma = \frac{\sum_{i=1}^{n_C} d_i}{T_C}, \quad \eta = \alpha\beta + \gamma$$

3.2. Beacon Length. A beacon needs to be transmitted entirely within a reception window of a receiving device for being received successfully. Each beacon has a certain transmission duration ω_i , and if the beacon transmission starts after the last ω_i time-units of a reception window (cf. after the start of the hatched area in Figure 2a)), it cannot be received successfully. Nevertheless, for simplicity of exposition, for now we assume that any overlap between a beacon and a reception window leads to a successful discovery. We further assume that all beacons have the same length ω and neglect the contribution of the transmission duration of the first successfully received beacon to the worst-case latency. We study the relaxation of these assumptions in Section 6.1 and 6.2.

4. DETERMINISTIC BEACON SEQUENCES

A device F can successfully discover another device E only if E sends a beacon during one of the reception windows of F . We refer to the other direction as E discovering F . In what follows, we first consider F discovering E only, and later generalize it towards mutual discovery.

On device E , let $B' = b_1, b_2, \dots$ be a subsequence of B_∞ . From here on, we will always assume that b_1 is the first beacon that is in range of a remote device F . This is because any prior beacons of B_∞ , when E is not within the range of F , are not relevant for ND. Further, let F run an infinite reception window sequence C_∞ . Though B_∞ and hence $B' \in B_\infty$ could be of infinite length, let us think of B' as a fixed-length sequence. This assumption is valid because in the case of a successful discovery, beacons that are sent thereafter are no longer relevant for the discovery procedure. Now recall that the reception windows of C_∞ are formed by concatenations of a finite sequence C and every instance of C has its own time origin, as defined by Definition 3.1 (cf. Figure 1a)). The first beacon b_1 in B' lies within a certain instance of C and has a certain (random) offset Φ_1 from the time origin of this instance of C . This is depicted in Figure 2b), which shows an infinite beacon sequence consisting of concatenations of $C = c_1, c_2, c_3, \dots$, of which one full instance is depicted. In addition, the figure contains the

last reception window c_3 of the preceding instance and the first reception window c_1 of the succeeding one. Further, three beacons b_0 , b_1 and b_2 are shown, of which only b_1 and b_2 are in range. Here, B' consists of b_1 , b_2 and some later beacons that are not shown in the figure. Beacon b_1 falls into the depicted instance of C and has an offset of Φ_1 time-units from its origin.

For some valuations of Φ_1 , at least one beacon of B' will coincide with a reception window of C_∞ . For other valuations of Φ_1 , there might be no beacon in B' that coincides with any reception window of C_∞ , irrespective of the length of B' . If an overlapping pair of a beacon and a reception window exists for **all** possible offsets Φ_1 , the tuple (B', C_∞) guarantees discovery within a bounded amount of time and hence realizes deterministic ND. We, in the following, formalize the properties that such a tuple (B', C_∞) needs to fulfill for guaranteeing discovery.

4.1. Coverage and Determinism. A tuple (C_∞, B') , along with Φ_1 , is depicted in Figure 2b). For a given (C_∞, B') , it is obvious that the offset Φ_1 , which is a measure of the shift between B' and C_∞ , solely determines whether a beacon in B' overlaps with a reception window in C_∞ or not. The time-duration after which such an overlap takes place, and hence the discovery latency, is also determined by Φ_1 . For which values of Φ_1 will beacon b_1 fall into one of the reception windows? Clearly, these are given by the set $\Omega_1 = \{[t_1, t_1 + d_1], [t_2, t_2 + d_2], \dots\}$ (cf. Figure 2b)). In other words, if Φ_1 lies within any interval belonging to Ω_1 , then b_1 is successfully received. Similarly, if Φ_2 is the offset of b_2 , then for Φ_2 belonging to any interval in Ω_1 , b_2 will be successfully received (see Figure 2b)).

Now, what are the offsets Φ_1 of b_1 , such that beacon b_2 is successfully received? These are given by the set $\Omega_2 = \{[t_1 - \lambda_1, t_1 + d_1 - \lambda_1], [t_2 - \lambda_1, t_2 + d_2 - \lambda_1], \dots\}$, where λ_1 is the time-distance between the beacons b_1 and b_2 , as already defined in Section 3 (see Figure 2b)). Therefore, Ω_2 is obtained by shifting all elements of Ω_1 by λ_1 time-units to the left. Similarly, $\Omega_3 = \{[t_1 - (\lambda_1 + \lambda_2), t_1 + d_1 - (\lambda_1 + \lambda_2)], [t_2 - (\lambda_1 + \lambda_2), t_2 + d_2 - (\lambda_1 + \lambda_2)], \dots\}$. Then Ω_k for $k = 3, 4, 5, \dots$ is similarly defined as

$$(3) \quad \Omega_k = \{[t_1 - \sum_{i=1}^{k-1} \lambda_i, t_1 + d_1 - \sum_{i=1}^{k-1} \lambda_i], [t_2 - \sum_{i=1}^{k-1} \lambda_i, t_2 + d_2 - \sum_{i=1}^{k-1} \lambda_i], \dots\}.$$

Now consider a beacon sequence $B' = b_1, \dots, b_m$ of length m . If Φ_1 belongs to any interval in $\Omega_1 \cup \Omega_2 \cup \dots \cup \Omega_m$, then one beacon from B' will be successfully received. We now extend this result and define a *coverage map*, which can be used to reason about valuations of the initial beacon offset Φ_1 that lead to successful discovery.

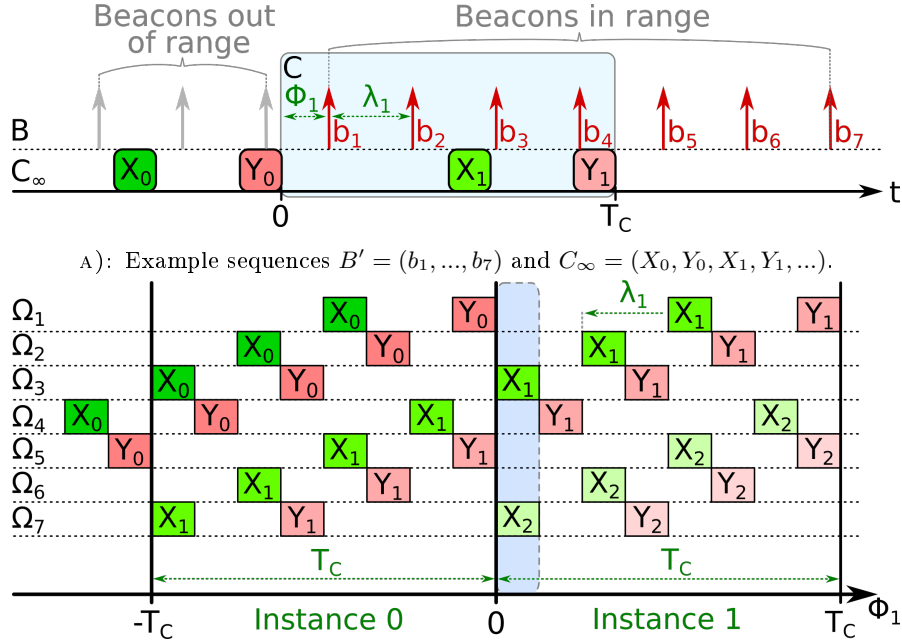
4.1.1. Coverage Maps. A coverage map is a formal representation of all offsets Φ_1 for which any beacon in B' overlaps with a reception window in C_∞ . It also allows for a graphical representation, from which several properties of the tuple (B', C_∞) can be easily understood.

Recall that C_∞ is a repeated concatenation of a sequence of reception windows C (i.e., $C_\infty = C C C \dots$). Now, we need to be able to specify specific instances of C within C_∞ . For this purpose, let us consider a simple example where C has two reception windows X and Y , and C_∞ is therefore given by $C_\infty = X Y X Y \dots$, and in order to distinguish between different instances of these reception windows, we will denote $C_\infty = X_0 Y_0 X_1 Y_1 X_2 Y_2 \dots$. The reception windows X_i and X_{i+1} , as well as Y_i and Y_{i+1} , are T_c time-units apart (see Figure 3a) and also Figure 1a)).

Figure 3a) shows a sequence of beacons $B' = b_1, \dots, b_7$ from a transmitting device. Below, two reception windows X_0, Y_0 from a receiving device are depicted, together with their periodic repetitions X_1, Y_1 , which are T_c time-units later. Again, $b_1 \in B'$ has a certain random offset Φ_1 from the origin of C . Figure 3b) shows the coverage map for the sequences in Figure 3a).

Definition 4.1 (Covered): An offset Φ_1 is *covered*, if at least one beacon in B' overlaps with any reception window in C_∞ for this offset.

Given the parameters of (B', C_∞) , the *construction* of a coverage map as in Figure 3b), is straightforward. We believe that the notion of such a coverage map and its use go beyond deriving latency bounds as done in this paper. It would also be useful for analyzing and optimizing various kinds of different ND protocols, including already known ones.



A): Example sequences $B' = (b_1, \dots, b_7)$ and $C_\infty = (X_0, Y_0, X_1, Y_1, \dots)$.
 B): Coverage map for these sequences drawn over two periods.
 FIGURE 3. Coverage maps.

From coverage maps, we can derive the following properties.

- **Beacon-to-beacon discovery latency l^* :** For a given offset Φ_1 , let $l^*(\Phi_1)$ be the latency measured from the transmission time of the first beacon that is in range, to the first time a beacon is successfully received. In Figure 3, $l^*(\Phi_1) = \tau_i - \tau_1 = \sum_{k=1}^{i-1} \lambda_k$, where i is the smallest row number in which Φ_1 is covered. For example, for an offset Φ_1 slightly above 0 (i.e., an offset within the highlighted frame in Figure 3b)), the beacon-to-beacon discovery latency will be $l^* = \tau_3 - \tau_1$, since b_3 is the earliest successful beacon for this offset.
- **Determinism:** By ensuring that all possible initial offsets are covered by at least one beacon, we can guarantee that B' is *deterministic* with respect to C_∞ (see next section for a formal definition of determinism).
- **Redundancy:** For certain valuations of Φ_1 , one can see in Figure 3b) that a beacon will be received by multiple reception windows. For example, for values of Φ_1 within the shaded frame, beacons b_3 and b_7 will be received by the windows X_1 and X_2 , respectively. By integrating over the length of all reception windows, for which such duplicate receptions happen, we can quantify the degree of redundancy of a tuple (B', C_∞) .

4.1.2. *Determinism.* Recall that protocols that can guarantee discovery for every possible initial offset are called *deterministic*. This is formalized below. In particular, we distinguish between a *beacon sequence* B' and a *protocol* (B_∞, C_∞) that can result in such a sequence.

Definition 4.2 (Deterministic ND Protocol): A beacon sequence B' is *deterministic* in conjunction with an infinite reception window sequence C_∞ , if all possible initial offsets Φ_1 are covered by the tuple (B', C_∞) . A ND protocol (B_∞, C_∞) is *deterministic*, if for all i , $B'_i = b_i, b_{i+1}, b_{i+2}, \dots$ is a deterministic beacon sequence.

Hence, deterministic ND protocols (B_∞, C_∞) always guarantee a bounded discovery latency, no matter when a beacon of B_∞ comes within the range of a receiving device.

Lemma 4.1: If a beacon sequence B' covers all offsets Φ_1 within $[0, T_C]$, then all possible valuations of Φ_1 are covered.

Proof. Let us assume that a certain range of offsets $[\Phi_x, \Phi_y]$, where $\Phi_x, \Phi_y \leq T_C$, is covered by a beacon b_i in conjunction with a certain reception window c_j . Since the pattern of reception windows repeats every T_C time-units, any $\Phi_1 \in [\Phi_x + T_C, \Phi_y + T_C]$ will result in b_i being received by the reception window c_{j+n_C} , which is T_C time-units after c_j . \square

Definition 4.3 (Redundant Sequences): If any offset Φ_1 within $[0, T_C]$ is covered by more than one beacon, then the tuple (B', C_∞) is *redundant*. Otherwise, (B', C_∞) is *disjoint*, since no intervals in the corresponding coverage map overlap.

For example, in Figure 3b), all offsets Φ_1 are covered and hence the corresponding tuple (B', C_∞) is deterministic. Further, since some offsets, e.g., the ones slightly above offset 0 (marked by the highlighted frame in Figure 3b)) are covered twice, it is also redundant.

4.1.3. *Coverage.* For a tuple (B', C_∞) , certain values of Φ_1 might be covered by multiple beacons, other values by exactly one beacon and yet others by no beacons. The notion of *coverage* quantifies how different values of $\Phi_1 \in [0, T_C]$ are covered. To understand this, recall that Ω_i is a set of intervals. Let us now consider those (full or partial) intervals of Ω_i that lie within $[0, T_C]$. The sum of the lengths of all such intervals for all Ω_i captures a notion of *coverage* that we formalize below.

Definition 4.4 (Coverage): Given a tuple (B', C_∞) , let a certain offset $\Phi_1 \in [0, T_C]$ be covered by k beacons, where $k \in \{0, 1, 2, \dots\}$. Let us define an auxiliary function $\Lambda^*(\Phi_1) = k$. Then, the *coverage* Λ is defined as

$$(4) \quad \Lambda = \int_0^{T_C} \Lambda^*(\Phi_1) d\Phi_1.$$

For example, in Figure 3b), if the lengths of X_i and Y_i are equal to unity and therefore $T_C = 8$, then $\Lambda = 14$. If $\Lambda < T_C$, a tuple (B', C_∞) cannot be deterministic, which implies that for certain values of Φ_1 , no bounded discovery latency can be guaranteed. If $\Lambda = T_C$, then (B', C_∞) can either be deterministic and disjoint, or else, it will be redundant and not deterministic. If $\Lambda > T_C$, then (B', C_∞) cannot be disjoint, and may or may not be deterministic.

4.2. **Minimum Coverage.** While Λ quantifies the coverage due to all beacons in B' , we now quantify the coverage induced by individual beacons.

Theorem 4.2 (Coverage per Beacon): Given a tuple (B', C_∞) , every beacon $b_i \in B'$ induces a coverage of exactly $\sum_{k=1}^{n_C} d_k$ time-units.

Proof. The first beacon b_1 in B' will cover exactly those time-units for which b_1 directly coincides with a reception window. The sum of such matching offsets is therefore $\sum_{k=1}^{n_C} d_k$ time-units. Every later beacon b_i will cover the same offsets shifted by the sum of beacon gaps $\sum_{k=1}^i \lambda_k$ to the left, which does not impact the amount of offsets covered. Since C_∞ is an infinite concatenation of a finite sequence C , for every covered offset that is shifted out of the considered range $[0, T_C]$, the same amount from a later period is shifted into that range, such that each beacon b_i covers exactly $\sum_{k=1}^{n_C} d_k$ time-units within $[0, T_C]$. \square

From the above, we are able to derive a minimum length of B' .

Theorem 4.3 (Beaconing Theorem): Given a tuple (B', C_∞) , the minimum number of beacons M a beacon sequence B' needs to consist of to guarantee deterministic discovery is:

$$(5) \quad M = \left\lceil \frac{T_C}{\sum_{k=1}^{n_C} d_k} \right\rceil$$

Proof. From Theorem 4.2 it follows that every beacon induces a coverage of $\Lambda = \sum_{k=1}^{n_C} d_k$. For deterministic discovery, the coverage Λ has to be at least T_C . Therefore, the number of beacons needed for deterministic ND must be at least $\lceil T_C/\Lambda \rceil$. \square




FIGURE 4. Partial sequences of an infinite beacon sequence.

It is worth mentioning that Theorem 4.3 is a necessary, but not sufficient condition for deterministic ND. The positioning of the beacons, along with their number, together determine whether or not a tuple (B', C_∞) is deterministic.

5. FUNDAMENTAL BOUNDS

In this section, we derive the lower bounds on the worst-case latency that a ND protocol could guarantee in different scenarios (e.g., symmetric or asymmetric discovery). In other words, given constraints like the duty-cycle, such a bound defines the best worst-case latency that any protocol could possibly realize. First, we consider the most simple case in which one device F runs an infinite reception window sequence $C_{F,\infty}$ without beaconing, whereas another device E only runs an infinite beacon sequence $B_{E,\infty}$ without ever listening to the channel. We refer to this as *unidirectional* beaconing.

5.1. Bound on Unidirectional Beaconing.

5.1.1. *The Coverage Bound.* Consider a tuple (B', C_∞) , where B' consists of M beacons and M is given by Theorem 4.3. Recall Theorem 4.3 and the subsequent discussion. If B' is disjoint and deterministic, then for every value of Φ_1 , there is exactly one beacon in B' that overlaps with a reception window in C_∞ . What are the beacon gaps λ_i using which such M beacons need to be spaced for minimizing the discovery latency?

The worst-case beacon-to-beacon discovery latency l^* , measured from the first beacon in range to the earliest successfully received one, is given by the sum of the $M - 1$ beacon gaps between these beacons. The first beacon in B' is the first beacon that was sent when the transmitter came within the range of the receiver. To measure the worst-case discovery latency L , time begins when the two devices come in range, which might be earlier than the time the first beacon in B' was sent. How much earlier? At most by the beacon gap that precedes B' . Recall that B' belongs to an infinite sequence B_∞ . Hence, the lowest worst-case latency is achieved if the sum of these M beacon gaps is minimized. At the same time, all offsets in $[0, T_C]$ need to be covered exactly once for ensuring determinism.

However, the following arguments rule out such M consecutive beacon gaps to be arbitrarily short. B_∞ has a transmission duty-cycle β , defined by the energy budget of the transmitter. Obviously, β determines the average beacon gap $\bar{\lambda}$. If the sum of a certain M consecutive beacon gaps becomes smaller than $M \cdot \bar{\lambda}$, then the sum of a different M consecutive beacon gaps within B_∞ needs to exceed $M \cdot \bar{\lambda}$ in order to respect average beacon gap of $\bar{\lambda}$ defined by β . Since any beacon in B_∞ could be the first beacon in range, the M beacons with the largest sum of beacon gaps determine the worst-case latency L . Hence, in an optimal B_∞ , every sum of M consecutive beacon gaps must be equal to $M \cdot \bar{\lambda}$. It is worth noting that this requirement does not necessarily require equal beacon gaps, because the above property has to hold for a specific value of M given by Theorem 4.3. This is formalized in Lemma 5.2.

To illustrate the above, consider the following example. Figure 4 shows a sequence $B' = b_1, \dots, b_7$. Here, let the minimum number M of beacons for deterministic ND be equal to 4 and let the partial sequences (b_1, \dots, b_4) , (b_2, \dots, b_5) , (b_4, \dots, b_7) be deterministic. Consider the sequence b_1, \dots, b_4 . Let us assume that b_4 would be sent somewhat earlier than depicted. Then, by decreasing λ_3 , the beacon gap λ_4 would increase accordingly, and

though the sequence b_1, \dots, b_4 would result in a shorter discovery latency for all possible offsets, the sequence b_4, \dots, b_7 would lead to a larger worst-case latency. The above observations are formalized below.

Theorem 5.1 (Coverage Bound): The lowest worst-case latency that can be guaranteed by a tuple (B_∞, C_∞) is:

$$(6) \quad L = \left\lceil \frac{T_C}{\sum_{i=1}^{n_C} d_i} \right\rceil \frac{\omega}{\beta}$$

Proof. Consider a sequence $B' = b_1, \dots, b_m$ with $m \gg M$. In B' , if any sum of M consecutive beacon gaps is less than $M \cdot \bar{\lambda}$, then the sum of a different M consecutive beacon gaps will exceed $M \cdot \bar{\lambda}$ and will define L . Since this is true for every m , it also holds for B_∞ . The mean beacon gap is given by $\bar{\lambda} = (\tau_m - \tau_1)/(m-1)$ and the worst-case latency by $L = M \cdot \bar{\lambda}$. Expressing the mean beacon gap by the duty-cycle for transmission (cf. Equation 1) and expanding M using Theorem 4.3 leads to Equation 6. \square

Lemma 5.2 (Repetitive Beacon Sequences): Given a repetitive C_∞ , every B_∞ that guarantees the lowest worst-case latency is repetitive, with a period of $m_B = M$ beacons or $T_B = M \cdot \frac{\omega}{\beta}$ time-units.

5.1.2. *Optimal Reception Window Sequences.* We know that in an optimal beacon sequence, the sum of every M consecutive beacon gaps is T_B . The corresponding reception window sequence must be such that all offsets in $[0, T_C]$ are covered by such a beacon sequence. While there can be multiple such C_∞ for a given B_∞ , the ones that are optimal must fulfill the following property.

Theorem 5.3 (Overlap Theorem): For a given duty-cycle, every C_∞ that minimizes the worst-case latency is given by the following equation.

$$(7) \quad T_C = k \cdot \sum_{i=1}^{n_C} d_i, \quad k \in \mathbb{N}$$

Proof. Let us assume that the length of T_C is equal to $k \cdot \sum_{i=1}^{n_C} d_i - \Delta$, where k is an integer and $\Delta \in [0, \sum_{i=1}^{n_C} d_i)$. Theorem 5.1 implies the same worst-case latency for all values of Δ , since the ceiling function in Equation 6 does not change L . With $T_C = k \cdot \sum_{i=1}^{n_C} d_i - \Delta$, the reception duty-cycle is given by (cf. Equation 2):

$$(8) \quad \gamma = \frac{\sum_{i=1}^{n_C} d_i}{k \cdot \sum_{i=1}^{n_C} d_i - \Delta}$$

From Equation 8 follows that the duty-cycle is minimized when $\Delta = 0$, and hence $T_C = k \cdot \sum_{i=1}^{n_C} d_i$. \square

The intuition behind Theorem 5.3 is that if Equation 8 is not satisfied, then T_C can be increased and therefore, the reception duty-cycle γ can be reduced without requiring any additional beacons to guarantee discovery with the same L . In other words, the coverage intrinsically induced if Equation 8 is not satisfied exceeds what is needed for determinism. By combining Theorem 5.1 and 5.3, we can derive a bound for unidirectional beaconing.

Theorem 5.4 (Fundamental Bound for Unidirectional Beaconing): Given a device E that runs an infinite beacon sequence $B_{E,\infty}$ with a duty-cycle of β_E and a device F that runs an infinite reception window sequence $C_{F,\infty}$ with a duty-cycle of γ_F , the minimum worst-case latency that can be guaranteed for F discovering E is as follows.

$$(9) \quad L = \left\lceil \frac{1}{\gamma_F} \right\rceil \frac{\omega}{\beta},$$

Clearly, optimal values of γ_F are of the form $1/k$, $k \in \mathbb{N}$ and other values of γ_F do not lead to an improved L compared to them.

Proof. By combining $T_C = k \cdot \sum_{k=1}^n d_k$ from Theorem 5.3 and Equation 1, we can write Equation 6 as follows.

$$(10) \quad L = \frac{T_C}{\sum_{i=1}^{n_C} d_i} \cdot \frac{\omega}{\beta} = \frac{\omega}{\beta \cdot \gamma}$$

This holds true for γ_F in the form of $1/k$, $k \in \mathbb{N}$. The proof for other duty-cycles follows from the above discussion. \square

5.2. Symmetric ND Protocols. In this section, we extend Theorem 5.4 towards bidirectional (i.e., device E discovers device F and vice-versa), symmetric (i.e., both devices E and F use the same duty-cycle η) ND. For achieving bidirectional discovery, every device runs both a beacon and a reception window sequence, and we assume that B_∞ and C_∞ can be designed such that both sequences on the same device never overlap with each other. We relax this assumption in Appendix B.

5.2.1. Bi-Directional Discovery. We can achieve bidirectional discovery by running the optimal sequences B_∞ and C_∞ we have identified for unidirectional beaconing on both devices simultaneously. The latency of each partial discovery procedure (viz., the discovery of E by F and of F by E) is bounded by Theorem 5.4. As a result, the worst-case latency for both partial discoveries being successful will also be bounded by Theorem 5.4. Since both devices transmit and receive, we can optimize the share between β and γ , which leads to the following bound.

Theorem 5.5 (Symmetric Bound for Bi-Directional ND Protocols): For a given duty-cycle η , no bi-directional symmetric ND protocol (i.e. every device runs the same tuple (B_∞, C_∞)) can guarantee a lower worst-case latency than the following.

$$(11) \quad L = \min \left(\underbrace{\left\lceil \frac{2}{\eta} \right\rceil^2 \cdot \frac{\omega\alpha}{\eta \left\lceil \frac{2}{\eta} \right\rceil - 1}}_{\text{A}}, \underbrace{\left\lfloor \frac{2}{\eta} \right\rfloor^2 \cdot \frac{\omega\alpha}{\eta \left\lfloor \frac{2}{\eta} \right\rfloor - 1}}_{\text{B}} \right)$$

Proof. Because of Theorem 5.3, optimal reception duty-cycles are given by $1/\gamma = k, k = 1, 2, 3, \dots$. By inserting $\eta = \alpha\beta + \gamma$ (cf. Definition 3.5) into Equation 11 and setting $1/\gamma = k$, we obtain

$$(12) \quad L = \frac{k^2\omega\alpha}{k\eta - 1}, k \in \mathbb{N}$$

We now have to find the value of k that minimizes L . Let us for now allow non-integer values of k in Equation 12. By forming the first and second derivative of Equation 12 by k , one can show that a local minimum of L exists for $k = 2/\eta$, which is a non-integer number for most values of η . By analyzing dL/dk , we can further show that Equation 12 is monotonically decreasing for values of $k < 2/\eta$ and monotonically increasing for values of $k > 2/\eta$. Hence, the only integer values of k that potentially minimize L are $\lceil 2/\eta \rceil$ and $\lfloor 2/\eta \rfloor$. Inserting $k = \lceil 2/\eta \rceil$ or $k = \lfloor 2/\eta \rfloor$ into Equation 10 and taking the minimum latency among both possibilities leads to Equation 11. \square

In fact, Theorem 5.5 also holds true for unidirectional beaconing, if the joint duty-cycle $\eta = \alpha \cdot \beta_E + \gamma_F$ of two devices E and F is to be optimized. Further, one can easily see that for small values of η , the ceiling function in Equation 11 only marginally affects the value of L , which can therefore be approximated by

$$(13) \quad L = \frac{4\alpha\omega}{\eta^2}.$$

Even when both devices E and F transmit as well as receive, it is possible to design *unidirectional* protocols in which only one of the two devices, E or F , can discover the other. Here, the beacons on both devices contribute to a joint notion of coverage, leading to a reduced latency bound compared to the case where both devices can discover each other mutually. A bound for this possibility is given in Appendix C.

5.2.2. Collision-Constrained Discovery. For achieving the bound given by Theorem 5.5, we have assumed that the beacons of multiple devices never collide. This assumption is reasonable for a pair of radios, in which collisions only rarely occur. However, as soon as more than two radios are carrying out the ND procedure simultaneously, collisions become inevitable and some of the discovery attempts fail. As a result, some devices might discover each other after the theoretical worst-case latency has passed, or, depending on the protocol design, might not discover each other at all. Therefore, it is often required to limit the channel utilization and hence collision rate, which leads to an increased worst-case latency bound.

In protocols with disjoint sequences (i.e., every Φ_1 is covered exactly once), every collision will lead to a failure of discovering within L . The collision probability is solely determined by the channel utilization β . We in this

section study the worst-case latency that can be achieved if both η and β (and hence the collision probability) are given. We in addition discuss possibilities to reduce the number of failed discoveries for a given collision probability in Section 8.1.1.

Consider a number of S senders, of which each occupies the channel by a time-fraction of β . The first beacon of an additional sender that starts transmitting (or comes into range) at any random point in time will face a collision probability of (cf. [54]):

$$(14) \quad P_c = 1 - e^{-2(S-1)\cdot\beta}$$

Once a beacon has collided, the repetitiveness of infinite beacon sequences (cf. Lemma 5.2) implies that the fraction of later beacons colliding with this device is predefined. Nevertheless, since all offsets between the two sequences occur with the same probability, the collision probability of every individual beacon is given by Equation 14. When constraining the channel utilization to a maximum value β_m that must never be exceeded, the following latency bound applies.

Theorem 5.6 (Bound for Symmetric ND with Constrained Channel Utilization): For a given upper bound on the channel utilization β_m , no symmetric ND protocol can guarantee a lower worst-case latency than the following.

$$(15) \quad L = \begin{cases} \min(\mathbb{A}, \mathbb{B}), & \text{if } \eta \leq \gamma_o + \alpha\beta_m \\ \left\lceil \frac{1}{\eta - \alpha\beta_m} \right\rceil \cdot \frac{\omega}{\beta_m}, & \text{if } \eta > \gamma_o + \alpha\beta_m \end{cases}$$

Here, \mathbb{A} and \mathbb{B} are given by Equation 11 and $\gamma_o = 1/\lceil 2/\eta \rceil$, if $\mathbb{A} \leq \mathbb{B}$, and $1/\lfloor 2/\eta \rfloor$, otherwise.

Proof. Given η , if the channel utilization that results from choosing the optimal value of γ (see proof of Theorem 5.5) does not exceed β_m , the bound given by Equation 11 remains unchanged. Otherwise, the bound is obtained from Equation 9 by eliminating γ using $\eta = \alpha\beta_m + \gamma$ (cf. Definition 3.5). \square

5.3. Asymmetric Discovery. So far, we have assumed that two devices E and F have the same duty-cycle, i.e., $\eta_E = \eta_F = \eta$. Next, we study the latencies of *asymmetric* protocols with $\eta_E \neq \eta_F$. We thereby assume that each device knows the duty-cycle of and hence the sequences on its opposite device. This is relevant e.g., when connecting a gadget with limited power supply to a smartphone using BLE. Here, different sequences on both devices that account for the different power budgets can be determined in the specification documents of the service offered by the gadget. The case of every device being allowed to choose its duty-cycle autonomously during runtime is also relevant. The possible degradation of the optimal performance for this case needs to be studied in further work. Whereas all our previously presented bounds are actually reachable by practical protocols, the bound we present for asymmetric ND can only be reached for tuples of duty-cycles (η_E, η_F) , for which $2/\eta_F$ and $2/\eta_E$ are integers. For other duty-cycles, the achievable performance will lie slightly below.

Theorem 5.7 (Bound for Asymmetric ND): Consider two devices E and F with duty-cycles η_E and η_F , where $\frac{2}{\eta_F}$ and $\frac{2}{\eta_E}$ are integers. The lowest worst-case latency for two-way discovery is as follows.

$$(16) \quad L = \frac{4\alpha\omega}{\eta_E\eta_F}$$

Proof. According to Theorem 5.4, if $1/\gamma_E$ and $1/\gamma_F$ are integers, the lowest worst-case one-way discovery latency L_F for device F discovering device E and the latency L_E for the reverse direction are as follows.

$$(17) \quad L_F = \frac{\omega}{\gamma_F \cdot \beta_E}, \quad L_E = \frac{\omega}{\gamma_E \cdot \beta_F}$$

The global worst-case latency for two-way discovery is given by $L = \max(L_E, L_F)$. Because of this, every optimal asymmetric ND protocol must fulfill $L_F = L_E$, since in cases of e.g., $L_F > L_E$, one could decrease the reception duty-cycle γ_F of device F and still achieve the same two-way discovery latency L . From $L_E = L_F$ and Equation 17 follows that $\gamma_F/\gamma_E = \beta_F/\beta_E = \text{const} = \mu$. By substituting β_E by β_F/μ in L_F (cf. Equation 17) and by substituting $\gamma_F = \eta_F - \alpha\beta_F$, we obtain:

$$(18) \quad L_F = \frac{\omega\mu}{(\eta_F - \alpha\beta_F)\beta_F}$$

By differentiating L_F by β_F , we can show that L_F is minimal for $\beta_F = \eta_F/2\alpha$ and hence $\gamma_F = \eta_F/2$. Similarly, L_E has a local minimum at $\beta_E = \eta_E/2\alpha$. We note that if $2/\eta_E$ and $2/\eta_F$ are integers, also $1/\gamma_E$ and $1/\gamma_F$ are integers. When re-substituting μ by β_F/β_E and replacing β_F and β_E by their optimal values, we obtain Equation 17. \square

6. RELAXATION OF ASSUMPTIONS

In Section 4, for the sake of ease of presentation, we have made multiple simplifying assumptions. In this Section, we relax all assumptions that have an impact on the discovery latency, study how the fundamental bounds are impacted by this and numerically evaluate the difference between the ideal and real bounds. In particular, we consider the bound for unidirectional beaconing from Theorem 5.4. We thereby consider only optimal reception duty-cycles γ , since other values of γ do not lead to any improvement of L (see Theorem 5.4).

6.1. Successful Reception of All Beacons. Throughout the paper, we have assumed that also beacons that only partially overlap with a reception window are received successfully. To account for the fact that beacons cannot be received if their transmissions start within the last ω time-units of each reception window (since they must entirely overlap with the window), we have to artificially shorten the actual length of each reception window d_k by one beacon transmission duration ω when computing discovery latencies, while still accounting for the full length of each reception window in computations of the duty-cycle. As a result, the coverage per beacon Λ in Equation 6 from Theorem 5.1 needs to be reduced by one beacon transmission duration ω for each reception window, such that a modified bound can be given as follows.

$$(19) \quad L = \left\lceil \frac{T_C}{\sum_{k=1}^{n_C} (d_k - \omega)} \right\rceil \frac{\omega}{\beta}$$

Clearly, this increases the worst-case latency that can be achieved for a given reception duty-cycle $\gamma = \sum_{k=1}^{n_C} d_k/T_C$. From this and Equation 19 follows that for a given reception duty-cycle γ , the increase of L becomes larger for higher numbers of reception windows n_C per period T_C . Hence, the tightest bound can be achieved for $n_C = 1$, for which the term $\sum_{k=1}^{n_C} d_k$ becomes d_1 . Using this and by restricting γ to optimal values and hence setting $T_C = k \cdot (d_1 - \omega)$ (cf. Theorem 5.3), we can write Equation 19 as:

$$(20) \quad L = \frac{T_C \omega}{T_C \beta \gamma - \beta \omega}$$

By examining the first derivative of Equation 20, one can show that L becomes smaller for growing values of T_C . However, L cannot become arbitrarily large. If $n_C = 1$ (i.e., one window per reception period), T_C cannot exceed L time-units because of the following reason. Consider a beacon that is sent at the very beginning of a reception window of another device. Let us assume that two devices come into range infinitesimally after this beacon has been sent. Since optimal beacon sequences do not contain more than one beacon within L time-units that overlap the same reception window, the next successful beacon will overlap with the subsequent reception window, which begins T_C time-units later. In other words, T_C must not exceed L time-units, since L would otherwise scale with T_C . For $n_C > 1$, it is $T_C \leq n_C L$. Hence, we can to substitute T_C by L in Equation 20. Solving this equation by L leads to the following bound:

$$(21) \quad L = \frac{\omega + \beta \omega}{\beta \gamma}$$

One can show that this bound is independent of the number of reception windows n_C per reception period.

6.2. Neglecting the First Successful Beacon. Throughout the paper, we have neglected the transmission duration of the first successfully received beacon. We can account for this by adding ω time-units to Equation 19. By forming the first and second derivative, we can show that the optimal share between transmission and reception is not influenced by this. When accounting for this beacon, all our presented bounds become by ω time-units longer (e.g., Theorem 5.5 becomes $L = 4\alpha\omega/\eta^2 + \omega$). Besides from this, there are no changes, since finding the optimal beaconing duty-cycle β is the only step that is potentially sensitive on adding ω to L .

6.3. Radio Overheads. Throughout this paper, we have assumed that the radios do not require any energy to switch from sleep mode to transmission or reception, and vice-versa. We now assume an overhead d_{oTx} to switch the radio from the sleep mode to transmission and back, and an overhead d_{oRx} to switch from the sleep mode to reception and back. These overheads can be regarded as effective durations of additional active time, i.e., as the actual durations that are needed to switch the radio's mode of operation, weighted by the quotient of the average power consumption during the switching phase over the power consumption for reception. For the sake of simplicity of exposition, we also assume the same overheads for switching directly between reception and transmission, without going to a sleep mode in between.

When accounting for these offsets, one can derive from Equation 2 that the duty-cycle for transmission β and for reception γ become:

$$(22) \quad \beta = \sum_{i=1}^{m_B} \frac{\omega_i + d_{oTx}}{\lambda_i}, \quad \gamma = \frac{\sum_{k=1}^{n_C} (d_k + d_{oRx})}{T_C}$$

With this, Equation 9, from which all other bounds are derived, becomes:

$$(23) \quad L = \left[\frac{1}{\gamma - \frac{n_C d_{oRx}}{T_C}} \right] \cdot \frac{\omega + d_{oTx}}{\beta}$$

We only consider optimal values of γ , which are given by $\gamma = 1/k + n_C d_{oRx}/T_C$. Large values of T_C minimize L , and hence, as explained in Section 6.1, we set $T_C = n_C \cdot L$. This leads to the following bound.

$$(24) \quad L = \frac{d_{oTx} + \omega + \beta \cdot d_{oRx}}{\beta \gamma}$$

6.4. Evaluation. In this section, we numerically evaluate the impact of the simplifying assumptions described above on the latency bound.

We assume a transmission duration of 32 μ s, which corresponds to a 4-byte beacon on a 1 MBit/s - radio used for e.g, BLE. We consider a range of duty-cycles β of the sender and γ of the receiver between 0.055 % and 5.55 %. This range of duty-cycles leads to a practically relevant range of discovery latencies from 0.1s to 100s for optimal protocols on ideal hardware platforms (i.e., for Equation 9). We only consider optimal values of γ . We further assume $\alpha = 1$.

Let L_i denote the ideal latency bound (viz., Equation 9) and L_r the latency bound with relaxed assumptions. As can be seen from Figure 5, in the considered range of duty-cycles, the relative deviation $(L_r - L_i)/L_i$ ranges between nearly 0 % to 6 %.




FIGURE 5. Relative difference between real and ideal bound on radios without switching overheads.




FIGURE 6. Slotted schedule proposed in [21]. Hatched bars depict beacons, smaller rectangles reception windows.

While Figure 5 provides a platform-independent comparison for any ideal 1 MBit/s radio, what performance can be achieved on existing hardware platforms? For a Nordic nRF51822 SOC [55], the switching overheads are approximately given by $d_{oRx} = d_{oTx} = 140 \mu\text{s}$. Within the considered range of duty-cycles, the relative deviation to the ideal bound ranges between 438 % and 467 %.

7. PREVIOUSLY KNOWN PROTOCOLS

In this section, we relate the worst-case performance of popular protocols and previously known bounds to the fundamental limits described in the previous section. Due to their relevance in practice, we consider only small duty-cycles η , for which the bound for symmetric protocols is given by Equation 13.

7.1. Worst-Case Bound of Slotted Protocols. As already described in Section 2, a worst-case number of slots within which discovery can be guaranteed is known for slotted protocols [20, 21]. The corresponding worst-case latency in terms of time is proportional to the slot length I , for which there is no known lower limit. In this section, we for the first time transform this worst-case number of slots into a latency bound and establish the relations to the fundamental bounds on ND presented in this paper. We will also address the bound presented in [10, 11], which has been claimed to be tighter than the bound in [20, 21].

7.1.1. Latency/Duty-Cycle Bound. According to [20, 21], no symmetric slotted protocol can guarantee discovery within T slots by using less than $k \geq \sqrt{T}$ active slots per T . The associated worst-case latency L is $T \cdot I$ time-units, which is directly proportional to the slot length I . We in the following derive a theoretical lower limit for I and hence for L .

Slotted protocols can only function properly if the beacon length ω is “at least one order of magnitude smaller than I ” [21]. If this requirement is not fulfilled, often a beacon might not overlap with a reception window even though the active slots of two devices overlap, as illustrated in Figure 6. Here, the slot length I in a slot design as proposed in [21] has been set to $2 \cdot \omega$. As can be seen, practically none of the offsets for which two active slots overlap would lead to a successful reception, since every beacon would only partially overlap with a reception window. If I would be increased, the fraction of successful offsets would gradually become larger. For achieving zero collisions independently of the slot length, let us assume a full duplex radio, which can both transmit and receive during the same points in time. Then, the theoretical limit on the slot length I becomes as low as one beacon transmission duration ω , which leads to the following duty-cycle:

$$(25) \quad \eta = \frac{k \cdot (I + \alpha\omega)}{T \cdot I} = \frac{k \cdot (I + \alpha\omega)}{L}$$

Since the limit from [20, 21] requires that $k \geq \sqrt{T} = \sqrt{L/I}$, with a slot length of $I = \omega$, Equation 25 leads to the following latency limit:

$$(26) \quad L \geq \frac{\omega(1 + 2\alpha + \alpha^2)}{\eta^2}$$

For $\alpha = 1$, this bound becomes $\frac{4\omega}{\eta^2}$ and hence identical to the fundamental bound for symmetric protocols given by Theorem 5.5. For all other values of α , this bound exceeds the one given by Theorem 5.5.

However, the assumption of full-duplex radios is not fulfilled by most wireless devices. Further, every wireless radio requires a turnaround time to switch from transmission to reception, during which the radio is unable

to receive any beacons. Even for recent radios, this time is large against the beacon transmission duration ω (e.g, for the nRF51822 radio [55], it lies around 140 μs , whereas beacons can be as short as 32 μs). Therefore, I will be orders of magnitude larger than ω , which linearly increases the worst-case latency slotted protocols can guarantee in practice. It is worth mentioning that this increase occurs in addition to the duty-cycle overhead induced by the turnaround times of the radio.

We now study the bound presented in [10, 11], which has been claimed to be lower in terms of slots than the one presented in [20, 21]. It is achieved by assuming two beacon transmissions per active slot ([20, 21] assumes only one), of which one beacon is sent slightly outside of the slot boundaries. By accounting for the two beacons per active slot, Equation 25 becomes $\eta = \frac{k \cdot (I + 2\alpha\omega)}{L}$, which leads to the following bound for the protocols proposed in [10, 11]:

$$(27) \quad L \geq \frac{\omega(\frac{1}{2} + 2\alpha + 2\alpha^2)}{\eta^2}$$

This bound becomes minimal for $\alpha = 1/2$, for which it is identical to the bound in Theorem 5.5. Hence, the bound in [10, 11] is lower in terms of slots than the bound in [20, 21], but identical or larger in terms of time.

7.1.2. Latency/Duty-Cycle/Channel Utilization Bound. All previously known bounds for slotted protocols are in the form of relations between the worst-case number of slots and the duty-cycle. The channel utilization, which is directly related to the beacon collision rate, has not been considered before. However, in slotted protocols, the channel utilization depends both on the number of active slots per period and on the slot length. For sufficiently large slot lengths, the turnaround times of the radio only play a negligible role. Further, the time for reception in each slot approaches nearly the whole slot length I . Hence, for $I \gg \omega$, we can compute the duty-cycle of slotted protocols as follows.

$$(28) \quad \beta = \frac{k\omega}{IT}, \quad \gamma = \frac{kI}{IT} = \frac{k}{T}, \quad \eta = \gamma + \alpha\beta$$

With the requirement of $k \geq \sqrt{T}$ from [20, 21], one can express the slot length I by the desired channel utilization β in Equation 28, which results in the following bound.

$$(29) \quad L \geq \frac{\omega}{\eta\beta - \alpha\beta^2}$$

From comparing Theorem 5.6 (cf. Equation 15) to Equation 29, it follows that if β_m lies below $\eta/2\alpha$, the worst-case latency a slotted protocol can guarantee with a channel-utilization of $\beta = \beta_m$ is identical to the corresponding fundamental bound (recall that we only consider optimal duty-cycles). For $\beta_m > \eta/2\alpha$, slotted protocols cannot reach the fundamental bound from Theorem 5.6. In practice, this means that slotted protocols can potentially perform optimally in busy networks with many devices discovering each other simultaneously, but cannot offer optimal performance in networks in which new devices join gradually and hence only a master node and the joining device need to carry out ND at the same time.

We in the following evaluate the popular protocols Disco [7], Searchlight-Striped [9], U-Connect [8] and diffcode-based protocols [20] and compare them to the fundamental bound given by Theorem 5.6. In Disco, active slots are repeated after every p_1 and p_2 slots, where p_1 and p_2 are coprimal numbers. The Chinese Remainder Theorem implies that there is a pair of overlapping slots among two devices every $p_1 \cdot p_2$ time-units. U-Connect also relies on coprimal numbers for achieving determinism. In contrast, Searchlight defines a period of T and a hyper-period of T^2 slots. The first slot of each period is active, whereas a second active slot per period systematically changes its position, until all possible positions have been probed. Diffcode-based solutions are built on the theory of block designs and hence guarantee a pair of overlapping slots among two devices with the minimum possible number of active slots per worst-case latency. More details on these protocols can be found in [56].

Slot length-dependent equations on the worst-case latency and duty-cycle of these protocols are available from the literature. When assuming sufficiently large slots and by expressing the slot length I by the channel utilization β similarly to Equation 28, one can derive the relations between the worst-case latency, duty-cycle

Protocol	$L(\beta, \eta)$
Diffcodes [20]	$\frac{\omega}{\eta\beta - \alpha\beta^2}$
Disco [7]	$\frac{8\omega}{\eta\beta - \alpha\beta^2}$
Searchlight-S [9]	$\frac{2\omega}{\eta\beta - \alpha\beta^2}$
U-Connect [8]	$\frac{(3\omega + \sqrt{\omega^2(8\eta - 8\alpha\beta + 9)})^2}{8\omega\beta\eta - 8\omega\alpha\beta^2}$

TABLE 1. Worst-case latencies of slotted protocols.

and channel utilization given in Table 1. Clearly, only Diffcode-based schedules reach the optimal performance in this metric, whereas all other ones perform below the optimum.

In summary, slotted protocols can perform optimal in the latency/duty-cycle/channel utilization metric, if the channel utilization remains low. In the latency/duty-cycle metric, however, higher required channel utilizations prevent slotted protocols from performing optimally.

7.2. Worst-Case Bound of Slotless Protocols. In slotted protocols, the number of beacons is always coupled to the number of reception phases. Slotless protocols are not subjected to this constraint. Can they reach optimal latency/duty-cycle relations? In [18], two parametrization schemes for slotted protocols, called $PI-0M$ and $PI-kM^+$, have been proposed, which have been claimed to provide the best latency/duty-cycle performance among all known slotless protocols. We therefore in the following relate their performance to the bounds presented in Section 5.

In such slotless protocols, beacons are sent periodically with a period T_B and the device listens to the channel for d time-units once per period T_C . When optimizing T_B , T_C and d , the worst-case latency is given by (cf. [18] for details):

$$(30) \quad L = \left(\left\lceil \frac{T_C - d + \omega}{T_B} \right\rceil + 1 \right) \cdot T_B + \omega$$

As for our bounds, we assume that 1) beacons that are sent within the last ω time-units of each reception phase are received successfully and 2) the transmission duration of the first successfully received beacon is neglected. Under these assumptions, we can set $\omega = 0$ and obtain the following worst-case latency of $PI-0M$:

$$(31) \quad L = \left(\left\lceil \frac{T_C - d}{T_B} \right\rceil + 1 \right) \cdot T_B$$

Further, [18] requires that $T_B = d$ and $T_C = (M + 1)d - \Delta$, $M \in \mathbb{R}$ and $\Delta \rightarrow 0$. This leads to a worst-case latency L of $\frac{\omega(M+1)^2}{\eta(M+1)-1}$ time-units. By forming the first and second derivative of L , one can find the optimal value of M , using which L becomes $4\omega\alpha/\eta^2$. This is identical to Theorem 5.5 and hence, under the assumptions described above, the $PI-0M$ scheme is optimal in the latency/duty-cycle metric. One can also show that under ideal assumptions, $PI-kM^+$ performs optimally, while it performs slightly below $PI-0M$ under relaxed assumptions. Which degradation of the latency bound of the $PI-0M$ scheme do these assumptions imply in practice? When assuming a beacon transmission duration of $\omega = 32 \mu\text{s}$ and a range of duty-cycles between 0.1% and 100% in steps of 0.1%, the normalized root mean square error between Equation 11 and the equations presented in [18] for $PI-0M$ is 1.24%.

8. CONCLUSION

In this section, we first describe open problems left for future research and then summarize the main results of this paper.

8.1. Open Problems.

8.1.1. *Problems On Fundamental Limits.* Regarding the future work on fundamental limits, there are two important problems left open. First, what is the lowest latency an asymmetric protocol can guarantee, if the duty-cycles of all devices are unknown? An what is the bound for asymmetric ND for duty-cycles for which $2/\eta$ is not an integer?

Second, the bounds derived so far are valid for a pair of devices discovering each other. For unidirectional beaconing, protocols in which 100 % of all discovery attempts are successful within L time-units can be realized in practice. For increasing numbers of devices discovering each other simultaneously, it is inevitable that their beacons will collide and hence, an increasing number of discovery attempts will fail. Therefore, generalized performance bounds for multi-device scenarios need to be derived. Such bounds are in the form of a function $L(\beta, \gamma, S, P_f)$, which needs to be interpreted as follows. For a given number of devices S with duty-cycles β and γ , in no ND protocol, a fraction of at least $1 - P_f$ of all discovery attempts will terminate successfully within less than L time-units. Clearly, for $S \rightarrow 1$ and $P_f \rightarrow 0$, this bound converges to L from Equation 9. The following two mechanisms determine the performance in multi-device scenarios.

1) Lowering the Channel Utilization: The rate of collisions directly correlates to the channel utilization β , as described by Equation 14. Hence, devices can reduce the failure probability P_f by reducing β , which will, however, negatively affect the discovery latencies achieved in the two-device case (cf. Equation 9).

2) Redundant Coverage: Optimality in the $L(\beta, \gamma)$ - metric for two devices implies that every initial offset is covered exactly once (cf. Theorems 4.3 and 5.3) and hence, every collision leads to a failed discovery. However, an ND protocol might cover multiple or all initial offsets more than once. Hence, for such offsets, more than one beacon would overlap with a reception window, and as long as one of them is not subjected to collisions, the discovery procedure will succeed. Moreover, it seems feasible to construct protocols that first cover every offset exactly once by a beacon sequence B' of length M . In addition, the same offsets are then covered again by concatenations of multiple instances of B' . In other words, such protocols would guarantee short latencies in the two-device case while performing potentially optimally also in multi-device scenarios.

The collision of a pair of beacons from two devices often induces an increased collision probability of subsequent pairs of beacons. For example, consider protocols in which beacons are sent with periodic intervals. Since all devices transmit with the same interval, a collision implies that all later beacons will also collide. To make protocols robust against failures due to collisions, a beacon schedule needs to fulfill the following property. Given any two beacons that both overlap with a reception window for the same offset Φ_1 , their individual collision probabilities should exhibit the lowest possible correlation. It is currently not clear which degree of such a decorrelation can be actually achieved. Further, measures for decorrelating collision probabilities might reduce the latency performance, because they could prevent beacons from being sent at their optimal points in time. Hence, not all initial offsets can be covered with the fewest possible number of beacons, making additional beacon transmissions necessary. Besides open questions on decorrelating collisions, for protocols being optimal in the multiple-device case, how many times should every initial offset be covered? These questions need to be studied further in order to derive agnostic bounds in the form of $L(\beta, \gamma, S, P_f)$.

8.1.2. *Problems in Protocol Design.* Our results also outline two important directions for the development of future ND protocols. First, there is no existing protocol which, for *every* duty-cycle and every required collision rate, could realize the optimal performance predicted by Theorem 5.6. Second, protocols that contain decorrelation mechanisms to make the collision of each beacon independent from the occurrence of previous collisions have not been studied thoroughly. Though BLE applies some random delay for scheduling its beacons [57], the optimal randomization technique to obtain the best trade-off between robustness and worst-case latency remains an open question.

8.2. Concluding Remarks. We have presented and proven the correctness of multiple fundamental bounds on the performance of deterministic ND protocols. In particular, we have presented bounds for unidirectional beaconing, for symmetric and for asymmetric bi-directional ND. Further, we have shown that in the latency/duty-cycle metric, only slotless protocols can reach optimal performance. However, if the channel utilization is constrained, both slotted and slotless protocols can perform optimally. We have also revealed new important open problems to be addressed by future research.

ACKNOWLEDGEMENTS

This work was partially supported by the German Research Foundation (DFG) under grant number CH918/5-1 - "Slotless Neighbor Discovery". We gratefully acknowledge Prof. Polly Huang for shepherding our paper at SIGCOMM.

REFERENCES

- [1] M. J. McGlynn and S. A. Borbash, "Birthday protocols for low energy deployment and flexible neighbor discovery in ad hoc wireless networks," in *ACM International Symposium on Mobile Ad Hoc Networking & Computing (MobiHoc)*, 2001.
- [2] R. Margolies, G. Grebla, T. Chen, D. Rubenstein, and G. Zussman, "Panda: Neighbor discovery on a power harvesting budget," *IEEE Journal on Selected Areas in Communications*, vol. 34, no. 12, pp. 3606–3619, 2016.
- [3] L. You, Z. Yuan, P. Yang, and G. Chen, "Aloha-like neighbor discovery in low-duty-cycle wireless sensor networks," in *IEEE Wireless Communications and Networking Conference (WCNC)*, 2011.
- [4] S. Vasudevan, D. Towsley, D. Goeckel, and R. Khalili, "Neighbor discovery in wireless networks and the coupon collector's problem," in *Annual International Conference on Mobile Computing and Networking (MobiCom)*, 2009.
- [5] C. Schurgers, V. Tsiatsis, S. Ganeriwal, and M. Srivastava, "Optimizing sensor networks in the energy-latency-density design space," *IEEE Transactions on Mobile Computing (TMC)*, vol. 1, no. 1, pp. 70–80, 2002.
- [6] Y. Tseng, C.-S. Hsu, and T.-Y. Hsieh, "Power-saving protocols for IEEE 802.11 based multi-hop ad hoc networks," in *IEEE Conference on Computer Communications (INFOCOM)*, 2002.
- [7] P. Dutta and D. Culler, "Practical asynchronous neighbor discovery and rendezvous for mobile sensing applications," in *ACM Conference on Embedded Network Sensor Systems (SenSys)*, 2008, pp. 71–84.
- [8] A. Kandhalu, K. Lakshmanan, and R. Rajkumar, "U-connect: A low-latency energy-efficient asynchronous neighbor discovery protocol," in *International Conference on Information Processing in Sensor Networks (IPSN)*, 2010, pp. 350–361.
- [9] M. Bakht, M. Trower, and R. Kravets, "Searchlight: Won't you be my neighbor?" in *Annual International Conference on Mobile Computing and Networking (MOBICOM)*, 2012, pp. 185–196.
- [10] T. Meng, F. Wu, and G. Chen, "On designing neighbor discovery protocols: A code-based approach," in *IEEE Conference on Computer Communications (INFOCOM)*, 2014, pp. 1689–1697.
- [11] —, "Code-based neighbor discovery protocols in mobile wireless networks," *IEEE/ACM Transactions on Networking (TON)*, vol. 24, no. 2, pp. 806–819, 2016.
- [12] L. Chen, R. Fan, L. Chen, M. Gerla, T. Wang, and X. Li, "On heterogeneous neighbor discovery in wireless sensor networks," in *IEEE Conference on Computer Communications (INFOCOM)*, 2015, pp. 693–701.
- [13] D. Zhang, T. He, Y. Liu, Y. Gu, F. Ye, R. K. Ganti, and H. Lei, "Acc: Generic on-demand accelerations for neighbor discovery in mobile applications," in *ACM Conference on Embedded Network Sensor Systems, (SenSys)*, 2012.
- [14] W. Sun, Z. Yang, W. Keyu, and L. Yunhao, "Hello: A generic flexible protocol for neighbor discovery," in *IEEE Conference on Computer Communications (INFOCOM)*, 2014, pp. 540–548.
- [15] Y. Qiu, S. Li, X. Xu, and Z. Li, "Talk more listen less: Energy-efficient neighbor discovery in wireless sensor networks," in *IEEE Conference on Computer Communications (INFOCOM)*, 2016, pp. 1–9.
- [16] C. Julien, C. Liu, A. L. Murphy, and G. P. Picco, "Blend: Practical continuous neighbor discovery for bluetooth low energy," in *ACM/IEEE International Conference on Information Processing in Sensor Networks (IPSN)*, 2017, pp. 105–116.
- [17] P. Kindt, D. Yunge, G. Reinert, and S. Chakraborty, "Griassdi: Mutually assisted slotless neighbor discovery," in *ACM/IEEE International Conference on Information Processing in Sensor Networks (IPSN)*, 2017, pp. 93–104.
- [18] P. Kindt, M. Saur, and S. Chakraborty, "Slotless protocols for fast and energy-efficient neighbor discovery," *CoRR*, vol. abs/1605.05614, 2016.
- [19] L. Wei, B. Zhou, X. Ma, D. Chen, J. Zhang, J. Peng, Q. Luo, L. Sun, D. Li, and L. Chen, "Lightning: A high-efficient neighbor discovery protocol for low duty cycle wsns," *IEEE Communications Letters*, vol. 20, no. 5, pp. 966–969, 2016.
- [20] R. Zheng, J. C. Hou, and L. Sha, "Asynchronous wakeup for ad hoc networks," in *ACM International Symposium on Mobile Ad Hoc Networking & Computing (MobiHoc)*, 2003.

- [21] R. Zheng, J. Hou, and L. Sha, "Optimal block design for asynchronous wake-up schedules and its applications in multihop wireless networks," *IEEE Transactions on Mobile Computing (TMC)*, vol. 5, no. 9, pp. 1228–1241, 2006.
- [22] T. Chen, J. Ghaderi, D. Rubenstein, and G. Zussman, "Maximizing broadcast throughput under ultra-low-power constraints," *IEEE/ACM Transactions on Networking (TON)*, vol. 26, no. 2, pp. 779–792, 2018.
- [23] A. Kandhalu, A. Xhafa, and S. Hosur, "Towards bounded-latency bluetooth low energy for in-vehicle network cable replacement," in *International Conference on Connected Vehicles and Expo (ICCVVE)*, 2013, pp. 635–640.
- [24] W. Lee, J. Youn, and T. Song, "Prime-number-assisted block-based neighbor discovery protocol in wireless sensor networks," *International Journal of Distributed Sensor Networks*, vol. 15, no. 1, p. 1550147719826240, 2019.
- [25] K. Kim, H. Roh, W. Lee, S. Lee, and D.-Z. Du, "PND: a p-persistent neighbor discovery protocol in wireless networks," *Wireless Communications and Mobile Computing*, vol. 13, no. 7, pp. 650–662, 2013.
- [26] S. Vasudevan, M. Adler, D. Goeckel, and D. Towsley, "Efficient algorithms for neighbor discovery in wireless networks," *IEEE/ACM Transactions on Networking (TON)*, vol. 21, no. 1, pp. 69–83, 2013.
- [27] A. Purohit, B. Priyantha, and J. Liu, "Wiflock: Collaborative group discovery and maintenance in mobile sensor networks," in *ACM/IEEE International Conference on Information Processing in Sensor Networks (IPSN)*, 2011.
- [28] S. Vasudevan and D. Kurose, J. and Towsley, "On neighbor discovery in wireless networks with directional antennas," in *IEEE Conference on Computer Communications (INFOCOM)*, vol. 4, 2005.
- [29] G. Jakkari, W. Luo, and S. V. Krishnamurthy, "An integrated neighbor discovery and MAC protocol for ad hoc networks using directional antennas," *IEEE Transactions on Wireless Communications (TWC)*, vol. 6, no. 3, pp. 1114–1024, 2007.
- [30] Z. Zhang and B. Li, "Neighbor discovery in mobile ad hoc self-configuring networks with directional antennas: algorithms and comparisons," *IEEE Transactions on Wireless Communications (TWC)*, vol. 7, no. 5, pp. 1540–1549, 2008.
- [31] N. Karowski, A. C. Viana, and A. Wolisz, "Optimized asynchronous multi-channel neighbor discovery," in *IEEE Conference on Computer Communications (INFOCOM)*, 2011.
- [32] W. Zeng, S. Vasudevan, X. Chen, B. Wang, A. Russell, and W. Wei, "Neighbor discovery in wireless networks with multipacket reception," in *Symposium on Mobile Ad Hoc Networking and Computing (MobiHoc)*, 2011.
- [33] S. Borbash, A. Ephremides, and M. J. McGlynn, "An asynchronous neighbor discovery algorithm for wireless sensor networks," *Ad Hoc Networks*, vol. 5, no. 7, pp. 998–1016, 2007.
- [34] K. Wang, X. Mao, and L. Y., "BlindDate: A neighbor discovery protocol," in *International Conference on Parallel Processing (ICPP)*, 2013, pp. 120–129.
- [35] Y. Zhang, K. Bian, L. Chen, P. Zhou, and X. Li, "Dynamic slot-length control for reducing neighbor discovery latency in wireless sensor networks," in *IEEE Global Communications Conference (GLOBECOM)*, 2017.
- [36] X. Guo, B. Chen, and M. Chan, "Analysis and design of low-duty protocol for smartphone neighbor discovery," *IEEE Transactions on Mobile Computing (TMC)*, vol. 16, no. 12, pp. 3294–3307, 2017.
- [37] D. Zhang, T. He, F. Ye, R. K. Ganti, and H. Lei, "Neighbor discovery and rendezvous maintenance with extended quorum systems for mobile applications," *IEEE Transactions on Mobile Computing (TMC)*, vol. 16, no. 7, pp. 1967–1980, 2017.
- [38] Z. Cao, Z. Gu, Y. Wang, and H. Cui, "Panacea: A low-latency, energy-efficient neighbor discovery protocol for wireless sensor networks," in *IEEE Wireless Communications and Networking Conference (WCNC)*, 2018.
- [39] H. Chen, W. Lou, Z. Wang, and F. Xia, "On achieving asynchronous energy-efficient neighbor discovery for mobile sensor networks," *IEEE Transactions on Emerging Topics in Computing (TETC)*, 2017.
- [40] L. Chen, K. Bian, and M. Zheng, "Never live without neighbors: From single- to multi-channel neighbor discovery for mobile sensing applications," *IEEE/ACM Transactions on Networking (TON)*, vol. 24, no. 5, pp. 3148–3161, 2016.
- [41] W. Lee, S. Choi, N. Kim, J. H. Youn, and D. Moore, "Block combination selection scheme for neighbor discovery protocol," in *International Conference on Communication Systems and Network Technologies (CSNT)*, 2015.
- [42] D. Yang, J. Shin, J. Kim, and G. H. Kim, "Opeed: Optimal energy-efficient neighbor discovery scheme in opportunistic networks," *Journal of Communications and Networks*, vol. 17, no. 1, pp. 34–39, 2015.
- [43] H. Wang, J. Ma, Y. Liu, W. Liu, and L. Wang, "Bi-directional probing for neighbor discovery," in *IEEE International Conference on Computational Science and Engineering (CSE)*, 2014.
- [44] A. Hess, E. Hytiä, and J. Ott, "Efficient neighbor discovery in mobile opportunistic networking using mobility awareness," in *International Conference on Communication Systems and Networks (COMSNETS)*, 2014.
- [45] M. Zhang, L. Zhang, P. Yang, and Y. Yan, "McDisc: A reliable neighbor discovery protocol in low duty cycle and multi-channel wireless networks," in *IEEE International Conference on Networking, Architecture and Storage (NAS)*, 2013.
- [46] B. Li, W. Feng, L. Zhang, and C. Spanos, "DEPEND: Density adaptive power efficient neighbor discovery for wearable body sensors," in *IEEE International Conference on Automation Science and Engineering (CASE)*, 2013.
- [47] R. Cohen and B. Kapchits, "Continuous neighbor discovery in asynchronous sensor networks," *IEEE/ACM Transactions on Networking (TON)*, vol. 19, no. 1, pp. 69–79, 2011.
- [48] D. Yang, J. Shin, J. Kim, and C. Kim, "An energy-optimal scheme for neighbor discovery in opportunistic networking," in *IEEE Consumer Communications and Networking Conference (CCNC)*, 2009.
- [49] S. T. S. Portal, "Bluetooth Low Energy (BLE) enabled devices market volume worldwide, from 2013 to 2020 (in million units)." 2018, www.statista.com/statistics/750569/worldwide-bluetooth-low-energy-device-market-volume.
- [50] Bluetooth SIG, "Find me profile specification," June 2011, revision V10r00, available via bluetooth.org.
- [51] P. Kindt, M. Saur, and S. Chakraborty, "Neighbor discovery latency in BLE-like protocols," *IEEE Transactions on Mobile Computing (TMC)*, vol. 17, no. 3, pp. 617–631, 2018.

- [52] M. Bradonjic, E. Kohler, and R. Ostrovsky, “Near-optimal radio use for wireless network synchronization,” *Theoretical Computer Science*, vol. 453, pp. 14–28, 2012.
- [53] L. Barenboim, S. Dolev, and R. Ostrovsky, “Deterministic and energy-optimal wireless synchronization,” *ACM Transactions on Sensor Networks (TOSN)*, vol. 11, no. 1, pp. 13:1–13:25, 2014.
- [54] N. Abramson, “The aloha system: Another alternative for computer communications,” in *ACM fall joint computer conference*, 1970.
- [55] Nordic Semiconductor ASA, “nRF51822 product spec. v3.1,” 2014, available via nordicsemi.com.
- [56] L. Chen and K. Bian, “Neighbor discovery in mobile sensing applications,” *Ad Hoc Networks*, vol. 48, no. C, pp. 38–52, 2016.
- [57] Bluetooth SIG, “Specification of the Bluetooth system 5.0,” December 2016, volume 0, available via bluetooth.org.

APPENDICES

APPENDIX A. NON-REPETITIVE RECEPTION WINDOW SEQUENCES

Throughout this paper, we have restricted our considerations to infinite length reception window sequences C_∞ that are given by concatenations of some finite sequence C . Though all currently known deterministic ND protocols are constructed accordingly, reception window sequences that continuously alter over time are also feasible. In what follows, we study such sequences and establish that all our presented bounds remain valid for them.

Let us consider an arbitrary pattern of reception windows of infinite length C_∞ . Such a C_∞ is characterized by its reception duty-cycle γ . As in Section 4, we consider a sequence B' that consists of those beacons that are sent after both devices have come into range. Obviously, the first beacon $b_1 \in B'$ is received successfully if it directly overlaps with one of the reception windows. The fraction of time-units at which a transmission of b_1 leads to a reception is therefore identical to γ . Another beacon that is sent by λ_1 time units later leads to additional points in time at which b_1 can be sent, such that one beacon out of b_1, b_2 is received successfully. These additional points in time lie λ_1 time-units earlier. Hence, like in Section 4, such points in time for later beacons are given by translating those of earlier ones to the left. If every point in time is covered by exactly one such translation, the tuple (B', C_∞) is disjoint and deterministic, and hence potentially optimal. This holds also true for cases in which C_∞ is not an infinite concatenation of the same C . The number of beacons M that need to be sent for guaranteeing deterministic discovery is therefore identical to the number of translations of the reception pattern C_∞ , such that every point in time overlaps with exactly one such translation. It is:

$$(32) \quad M = \left\lceil \frac{1}{\gamma} \right\rceil$$

This is identical to Theorem 4.3, and hence all bounds remain unchanged.

APPENDIX B. IMPLICATIONS OF SAME SEQUENCES ON BOTH DEVICES

Throughout the paper, we have assumed that C_∞ does not impose any constraints on scheduling the beacons in B_∞ on the same device. In this section, we study the relaxation of this assumption.

B.1. Symmetric Sequences. We first study the case in which both devices E and F run the same tuple of sequences (B_∞, C_∞) . Here, B_∞ is designed such that a beacon overlap with C_∞ is guaranteed for all initial offsets. Hence, not only an overlap of a beacon of B_F with $C_{E,\infty}$ is guaranteed, but also an overlap of a beacon of B_E with $C_{E,\infty}$. Such an overlap implies that the affected reception window needs to be interrupted for a certain amount of time.

For an ideal radio (i.e., a radio that does not require any time to switch from reception to transmission and vice-versa, see Section 6.3), this amount of time is identical to one beacon transmission duration ω . A beacon sent by another device within this period of time would collide and therefore would not be received successfully, even if the radio was able to receive and transmit simultaneously.

However, a real-world radio needs a certain amount of time d_{oTxRx} to switch from transmission to reception and an overhead d_{oRxTx} to switch from reception to transmission, during which no communication can be carried out. We in the following analyze the impact of this. Towards this, we next compute the time-fraction of all reception windows in C_∞ , during which the radio is unable to receive.

Since an optimal tuple of sequences (C_∞, B') is designed such that every initial offset is covered exactly once, exactly one beacon of B' will overlap with a reception window for every possible initial offset. For every such overlap, the radio is unable to receive incoming beacons for $d_{oTxRx} + d_{oRxTx} + \omega$ time-units within the affected reception windows.

In a tuple B_∞, C_∞ , how frequent do such overlaps occur and which fraction of the total reception time is “blocked” by them? In optimal protocols, exactly one beacon overlaps with a reception window per worst-case latency L (cf. Section 5.1). From Theorem 10 follows that for optimal values of γ , $L = T_C \cdot 1/(\sum_{i=1}^{n_C} c_i) \cdot \omega/\beta$, and hence L is always divisible by T_C . In every instance of T_C , there are $\sum_{i=1}^{n_C} d_i$ time-units during which the radio is scanning, and therefore, the radio spends $L/T_C \cdot \sum_{i=1}^{n_C} d_i = \omega/\beta$ time-units per worst-case latency L for scanning. The probability of failed discoveries is identical to the fraction of “blocked” time per L , which leads to the following equation.

$$(33) \quad P_{fail} = \frac{\beta}{\omega} \cdot (d_{oTxRx} + d_{oRxTx} + \omega)$$

In this equation, we assume that the amount of time during which the radio is “blocked” per beacon that overlaps with a reception window of the same device is always identical to $d_{oTxRx} + d_{oRxTx} + \omega$ time-units. We in the following prove this assumption.

Recall from Section 4.1 that every beacon of a deterministic Sequence B' , in conjunction with a reception window from C_∞ of a remote device, leads to a certain contiguous range of covered offsets, which we in the following call a *coverage image*. If the initial offset Φ_1 lies within one of these coverage images, B' is received successfully. Figure 7 exemplifies a coverage map of a non-redundant and deterministic (and hence potentially optimal) ND protocol. Here, $C \in C_\infty$ consists of only one reception window and hence, there is one coverage image per beacon. Recall that if a remote device sends a beacon during the last ω time-units of every scan window, it is not received successfully (cf. Section 6.1). We can therefore subdivide every coverage image of an optimal protocol into the following three parts \mathcal{A} , \mathcal{B} and \mathcal{C} (cf. Figure 7).

- Part \mathcal{C} has a length of ω time-units, and a beacon of the remote device that falls into this part will not be received successfully. Therefore, such Parts \mathcal{C} do not contribute to the overall coverage.
- To nevertheless ensure discovery if a beacon falls into such a Part \mathcal{C} of a coverage image, each Part \mathcal{C} is also covered by the Part \mathcal{A} of another coverage image, which also has a length of ω time-units.
- The remaining part \mathcal{B} is disjoint, i.e., no part of any other coverage image overlaps with it.

On a device E , we know that exactly one beacon of $B_{E,\infty}$ will overlap with at least one reception window of $C_{E,\infty}$ per L , which effectively interrupts or shortens the affected scan window. Such an overlap could happen in one of the following three ways.

- (1) The overlapping beacon falls into Part \mathcal{B} , such that a contiguous duration of $d_{oTxRx} + d_{oRxTx} + \omega$ time-units is blocked (e.g., it falls into the center of Part \mathcal{B}).
- (2) The overlapping beacon falls into the beginning (e.g., Part \mathcal{A}) of the scan window. Therefore, the “blocked” amount of time would also overlap with the neighboring Part \mathcal{B} (cf. Figure 7). Hence, the amount of occupied scanning time is equal to $d_{oTxRx} + d_{oRxTx} + \omega$ also for this situation.
- (3) The same holds true for a beacon falling into the end of the scan window (e.g., into Part \mathcal{C}), where parts of the “blocked” amount of time overlap with a Part \mathcal{A} and possibly also \mathcal{B} of another scan window.

Hence, in all three cases, the amount of “blocked” time is $d_{oTxRx} + d_{oRxTx} + \omega$.




FIGURE 7. Coverage map of a deterministic beacon sequence B'_F in conjunction with a certain $C_{E,\infty}$. The offsets covered by any reception window are composed by a Part \mathcal{A} that overlaps with the last ω time-units of another reception window, a Part \mathcal{B} that is disjoint and a Part \mathcal{C} , during which an incoming beacon is not successfully received.

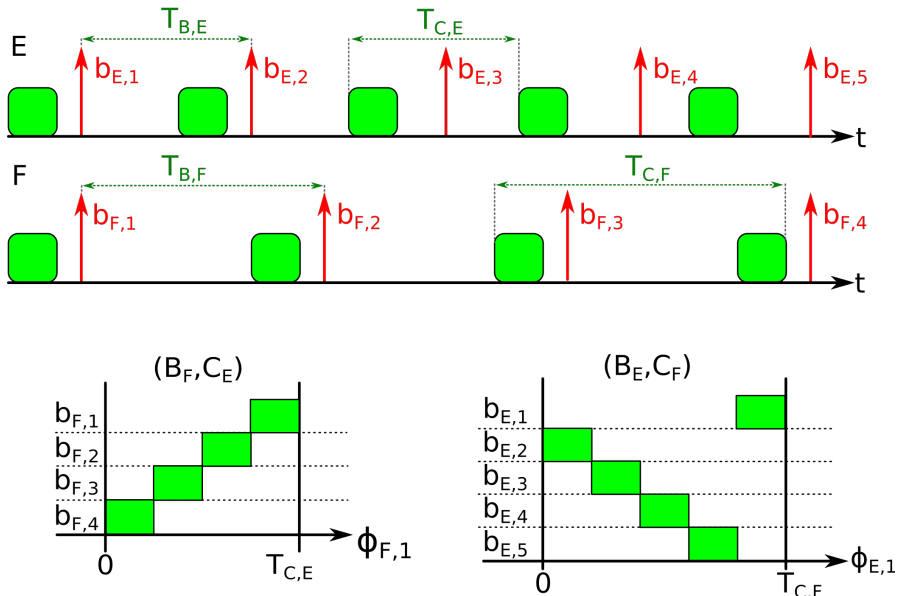


FIGURE 8. Asymmetric sequences and their coverage maps.

B.2. Asymmetric Sequences. For asymmetric discovery (i.e., both devices have different duty-cycles), a quadruple of beacon- and reception window sequences can be designed such that B_∞ and C_∞ on the same device never overlap, while allowing for optimal (i.e., disjoint coverage) and deterministic two-way discovery between the two devices. Figure 8 depicts a pair of tuples $(B_{F,\infty}, C_{F,\infty})$ and $(B_{E,\infty}, C_{E,\infty})$ along with the corresponding coverage maps. As can be seen, $(B_{F,\infty}, C_{E,\infty})$ and $(B_{E,\infty}, C_{F,\infty})$ realize disjoint and deterministic discovery, while the sequences on the same device never overlap.

APPENDIX C. MUTUAL EXCLUSIVE UNIDIRECTIONAL DISCOVERY

In Section 5.1, we have studied unidirectional discovery in the sense that one device F could discover E without E discovering F . However, it is also possible to design the tuple (B_∞, C_∞) on each device such that either device E or F can directly discover its opposite, which we study in this section.




FIGURE 9. Correlated offsets $\Phi_{F,1}$ and $\Phi_{E,1}$ in the sequences of two devices E and F .

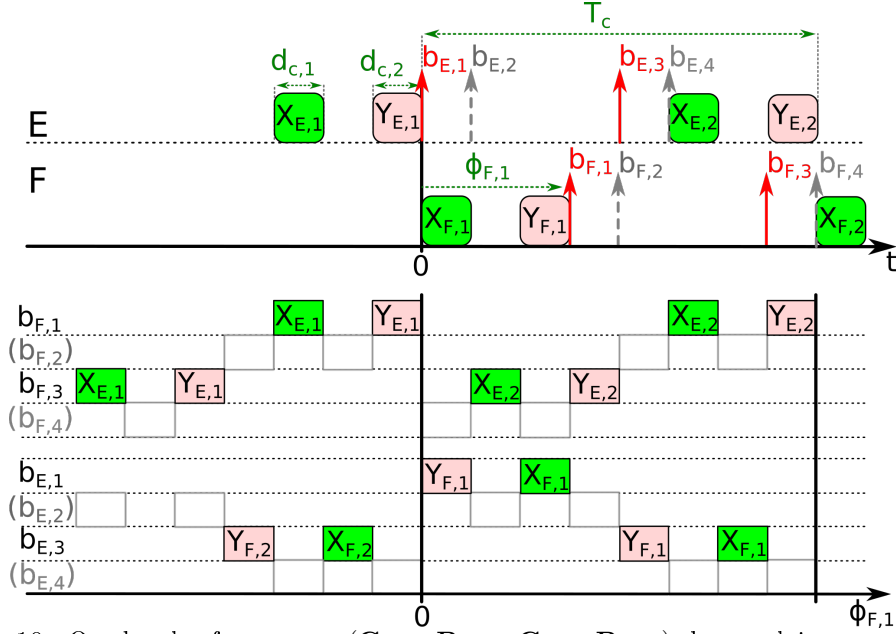


FIGURE 10. Quadruple of sequences $(\mathbf{C}_{E,\infty}, \mathbf{B}_{E,\infty}, \mathbf{C}_{F,\infty}, \mathbf{B}_{F,\infty})$ that exploits temporal correlations.

This form of unidirectional discovery is realized using beacon sequences $B \in \mathcal{B}_\infty$, in which the beacons on one device are sent with a fixed temporal relation to the reception windows on the same device. For example, let beacon $b_{F,1}$ on device F be sent by ζ time-units after reception window $X_{F,1}$, as depicted in Figure 9. Further, let such a relation exist on both devices and in every period T_C of the reception window sequence. As previously explained $b_{F,1}$ has a random offset of $\Phi_{F,1}$ time-units from the coordinate origin of device E . The temporal correlation between B_∞ and C_∞ on every device implies that the offset $\Phi_{E,1}$ beacon $b_{E,1}$ has from the coordinate origin of device F is fully determined by $\Phi_{F,1}$ (cf. Figure 9). It is:

$$(34) \quad \Phi_{E,1} = \zeta + (\zeta - \Phi_{F,1}) = 2 \cdot \zeta - \Phi_{F,1}$$

By exploiting this temporal relation, a quadruple of sequences $(C_{F,\infty}, B_{E,\infty}, C_{E,\infty}, B_{F,\infty})$ can guarantee deterministic one-way discovery, even if the pair $(C_{E,\infty}, B_{F,\infty})$ only covers half of the offsets $\Phi_{F,1} \in [0, T_C]$, by having the pair $(C_{F,\infty}, B_{E,\infty})$ covering the remaining ones. Thereby, the number of beacons that need to be sent per device for guaranteeing one-way discovery can be halved. The upper part of Figure 10 depicts the beacons (arrows) and reception windows (rectangles) of two devices E and F . On each device, the reception windows and beacons have a fixed temporal relation, whereas beacon $b_{F,1}$ has a random offset $\Phi_{F,1}$ from the coordinate origin of device E . Dashed arrows depict beacons that would need to be sent if every device would have to cover all offsets in the entire period T_C on its own. When exploiting temporal correlations between B_∞ and C_∞ on the same device, these beacons can be omitted without increasing the one-way worst-case latency. The lower part of Figure 10 depicts the coverage map of the beacons $b_{F,1}, \dots, b_{F,4}$ of device F and $b_{E,1}, \dots, b_{E,4}$ of device E . This coverage map represents all offsets $\Phi_{F,1}$, for which either a beacon from device F overlaps with a reception window of device E or a beacon from device E overlaps with a reception window of device F . Covered offsets of omitted beacons have been left white. As can be seen, every possible initial offset $\Phi_{F,1}$ is covered by either a beacon of $B_{F,\infty}$ falling into a reception window of $C_{E,\infty}$, or a beacon of $B_{E,\infty}$ falling into a reception window of $C_{F,\infty}$, and hence the number of beacons per device is halved compared to direct bi-directional discovery. This leads to the following latency bound, which is lower than the one given by Theorem 5.5. Since there are no further possibilities to improve pairwise discovery, this is also the tightest fundamental bound for all pairwise deterministic ND protocols.

Theorem C.1: The lowest worst-case latency a pair of devices can guarantee for mutual exclusive one-way discovery (i.e., either of both devices can discover its opposite one) is given by:

$$(35) \quad L = \min \left(\left\lceil \frac{1}{\eta} \right\rceil^2 \cdot \frac{\omega\alpha}{\eta \cdot \lceil 1/\eta \rceil - 1/2}, \left\lfloor \frac{1}{\eta} \right\rfloor^2 \cdot \frac{\omega\alpha}{\eta \cdot \lfloor 1/\eta \rfloor - 1/2} \right)$$

Proof. For a given set of offsets Ω_F covered by $B_F \in B_{F,\infty}$ on device E , Equation 34 defines a set of offsets Ω_E that are automatically covered by $B_E \in B_{E,\infty}$ on device F , and vice-versa. If Ω_F and Ω_E are disjoint, the amount of offsets contained in $\Omega_F \cup \Omega_E$ must sum up to T_C time-units for guaranteeing one-way discovery. Hence, the beacon sequence on every device needs to cover only $1/2 \cdot T_C$ time-units to guarantee one-way determinism, and Equation 6 becomes:

$$(36) \quad L = \left\lceil \frac{T_C}{2 \cdot \sum_{k=1}^{n_C} d_k} \right\rceil \frac{\omega}{\beta}$$

The rest of this proof is identical to the one for direct symmetric discovery (cf. Theorem 5.5). \square

Theorem C.1 is valid for one-way discovery (i.e., device E discovers device F or vice-versa). An indirect reverse discovery can be realized as follows. Each device transmits its next point in time at which it listens to the channel along with its beacons. The receiving device then schedules an additional beacon at the received point in time. This technique is called *mutual assistance* [17], and is actually a form of synchronous connectivity. Here, the latency for two-way discovery will be increased by the maximum temporal distance between any beacon and its succeeding reception window on the same device. An upper bound for this penalty for two-way discovery is T_C time units, which can be reduced significantly in sequences with more than one reception window per period T_C .

APPENDIX D. TABLE OF SYMBOLS

α	Ratio of the power spent for transmission over the power spent for reception
β	Duty-cycle for transmission, which is equivalent to the channel utilization
β_m	Specified maximum channel utilization
η	Duty-cycle
γ	Duty-cycle for reception
Λ	Coverage of a beacon sequence B' given an infinite reception window sequence C_∞

$\Lambda^*(\Phi_1)$	Number of beacons that cover the offset Φ_1
λ_i	Gap between beacon i and beacon $i + 1$
μ	Constant ratio of the reception or beaconing duty-cycles of two devices
Ω_i	Set of offsets Φ_1 covered by beacon i
ω_i	Transmission duration of beacon i
Φ_i	Offset of the i 'th beacon of a beacon sequence from the coordinate offset in C
τ_i	Time at which beacon i is sent
ζ	Fixed temporal distance of a certain beacon from a certain reception window on the same device in every period T_C
B/B_∞	Finite/infinite beacon sequence
b_i	Beacon i
C/C_∞	Finite/infinite reception window sequence
c_i	Reception window i
d_i	Time duration of reception window i
d_{oRxTx}	Effective additional active time for switching from reception to transmission
d_{oRx}	Effective additional active time for switching from sleep to reception and vice-versa
d_{oTxRx}	Effective additional active time for switching from transmission to reception
d_{oTx}	Effective additional active time for switching from sleep to transmission and vice-versa
I	Slot length in a slotted protocol
L	Worst-case latency
l^*	Beacon-to-beacon latency: Worst-case latency measured from the first beacon in range to the last, successfully received one.
L_i	Ideal worst-case discovery latency
L_r	Worst-case discovery latency when relaxing all simplifying assumptions
M	Minimum number of beacons needed to deterministically match an infinite reception window sequence C_∞
m_B	Period of a repetitive beacon sequence (in terms of number of beacons)
n_C	Number of reception windows contained in a finite length reception window sequence C , whose concatenations form an infinite sequence C_∞
P_c	Collision probability
P_f	Probability of a failed discovery
P_{Tx}, P_{Rx}	Power consumption of a radio for transmission or reception, respectively
S	Number of transmitting devices
T_B	Time-period of a repetitive, infinite beacon sequence
T_C	Time between the ends of two consecutive instances of the finite reception window sequence C , whose concatenations form an infinite sequence C_∞
t_i	Point in time the reception window i begins at

APPENDIX E. LIST OF ACRONYMS

ND: neighbor discovery
MANET: mobile ad-hoc network
BLE: Bluetooth Low Energy
PI: periodic interval