Modes of inter-network interaction in beacon-enabled IEEE 802.15.4 networks

Laura Marie Feeney*, Michael Frey†, Viktoria Fodor‡ and Mesut Günes§
*Swedish Institute of Computer Science, Email: lmfeeney@sics.se
†Department of Computer Science, Humboldt-Universität zu Berlin, Email: frey@informatik.hu-berlin.de
‡School of Electrical Engineering, KTH Royal Institute of Technology, Email: vjfodor@kth.se
†Institute of Computer Science, University of Münster, Email: mesut.guenes@uni-muenster.de

***Abstract*—Future growth in Internet-of-Things applications will lead to an operating environment with many *independent* networks operating in the same location. Co-existence and inter-network interference in IEEE 802.15.4 networks is particularly interesting because of the popularity of the radio and limited number of "good" channels, making it likely that there will be nearby networks operating on the same channel.**

This paper presents an in-depth simulation study of inter-network interaction between beacon-enabled IEEE 802.15.4 networks using contention-based slotted CSMA and contention-free GTS allocations. We use simple scenarios intended to highlight the underlying modes of interaction between interfering networks. Our results reveal several complex behaviors, including episodes of large, slow oscillations in throughput and long periods of disconnection interspersed with short bursts of high throughput. These results have practical implications for the design and evaluation of protocols and applications for use in future IoT environments.

I. Introduction

This paper presents an in-depth simulation study of the performance of independent IEEE 802.15.4 PAN networks that are operating in the same location and using the same channel. Our experiments are intended to characterize modes of interaction between networks using contention-based slotted CSMA and contention-free slot allocation (GTS) communication. The results demonstrate a significant impact on performance.

Our interest in the problem of inter-network interaction is driven by the growth of wireless sensor networks (WSN) and the Internet-of-Things (IoT). Especially in dense urban and residential environments, there will be many different kinds of networks and applications, belonging to many different owners, operating in the same physical location. As a result, the future IoT operating environment will be characterized by the presence of *many independent, co-located networks* sharing unlicensed spectrum.

IEEE 802.15.4 is an important case study for this scenario for three reasons: One, it is widely used for IoT applications. Two, only a limited number of IEEE 802.15.4 channels avoid cross-technology interference from IEEE 802.11 – four in the best case and only one in the worst. Even if there are only a few co-located networks, it is not improbable that two or more of them will share a channel. Three, IEEE 802.15.4 is the basis for a wide range of protocols, such as ContikiMAC [1], ZigBee [2], 6LowPAN [3], and WirelessHART [4]. Although these protocols share a common radio technology and PHY layer, they have very diverse power saving and channel access behaviors. It is unclear how such networks will interact in practical deployment scenarios. This paper is a first step toward answering this question.

The contributions of the paper are as follows: We identify and explain a number of surprisingly complex behaviors, even in the case of just two interacting networks. Our results demonstrate three key modes of interaction that occur over timescales of hours or minutes: large oscillations between high and low throughput; slow deterioration and recovery; and extended periods of loss interspersed with bursts of high throughput (and vice versa). We also show that significantly increasing the duty cycle (and energy consumption) yields only a moderate performance improvement. More generally, we argue that a better understanding of these interactions is needed to inform the design of more resilient protocols and applications.

The rest of the paper is organized as follows: Sections II and III give a brief overview of IEEE 802.15.4 and present the experiment design. Section IV discusses the simulation results. Sections V and VI present related work and conclusions.

II. Background

IEEE 802.15.4 is the primary standard for low-rate personal area wireless networks (PAN): Common chipsets for the 2.4 GHz unlicensed bands provide 250 kbps raw data rate and indoor communication ranges of 10-20 m. The standard defines both synchronous (beacon-enabled) and asynchronous (non beacon-enabled) operating modes, as well as star, cluster-tree, and pure peer-to-peer topologies.

In this work, we consider the fundamental beacon-enabled star topology, which consists of a central PAN coordinator and a set of associated devices. The coordinator periodically transmits a beacon frame announcing its presence. A device that receives the beacon can request to associate with the coordinator, remaining associated as long as it continues to regularly receive beacons.

The beacon synchronizes the network and defines the superframe structure (Fig. 1). The interval between beacons defines the length of the superframe, which is divided into an active and an inactive period. During the inactive period, devices

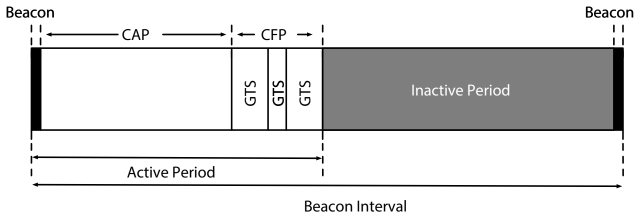

Fig. 1: The structure of the IEEE 802.15.4 superframe consisting of an active and inactive period with three exemplary guaranteed time slot (GTS) allocations.

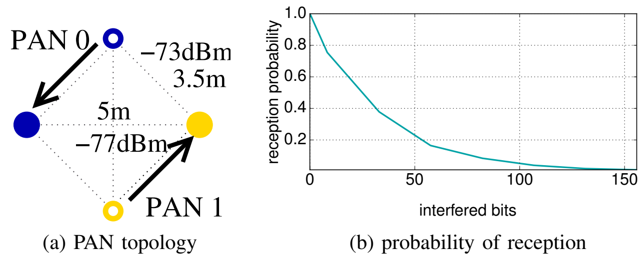

Fig. 2: PAN: Topology, RSS, and probability of reception under interference (sender at 3.5 m, interferer at 5 m).

can switch to a low-power sleep mode, waking up in time to receive the beacon at the beginning of the next active period.

The active period begins with a mandatory contention access period (CAP). During the CAP, senders contend for the channel using slotted CSMA, which is used for all management frames, as well as for data. Part of the active period can optionally be used as a contention-free period (CFP), from which devices can request a guaranteed time slot (GTS) allocation from the coordinator. Since GTS allocations are exclusive, devices send without sensing the channel or deferring during their GTS allocation. Because the beacon contains the allocation information for the following active period, a device is not allowed to send if it does not receive the beacon. If several consecutive beacons are lost, the device must reinitiate the association process.

III. Experiment design

Our experiments focus on the underlying modes of interaction between co-located PAN networks using contention-based slotted CSMA and contention-free GTS allocation. To make these interactions clearly visible, we use a stylized simulation scenario based on two PANs, each having just one associated device. This allows us to minimize the effects of in-network contention, as well as unfairness and variations in interference due to varying link quality.

Each interaction scenario is characterized by the time offset between the active periods of the two PANs, as defined by the offset betwee their beacon transmissions. We vary the beacon offset with very fine granularity, since arbitrary offsets can occur between two independent PANs that are within interference range of each other. More importantly, the networks will experience *all* possible beacon offsets over time, due to cumulative clock drift between their PAN coordinators. We also examine the relation between energy consumption and performance, by varying the duration of the active period within each beacon interval.

We base our investigations on simulation, as it provides complete control and visibility into system behavior. In particular, we require fine-grain *control* ($< 500\ \mu$s) over timing differences between PAN coordinators, which would be difficult to maintain without specialized testbed hardware. Moreover, our interest is in dense IoT environments, such as networks located in the same room, where the dominant effects will be high signal strength and strong interference. This reduces the urgency of detailed modeling of complex propagation, detection, interference, and frame reception behaviors, which tend to be a weakness of simulation compared to a testbed.

The rest of this section describes the simulation scenarios and the Castalia simulation environment; simulation parameters are shown in Table I.

A. Topology and Traffic

We simulate two PANs, each having a coordinator and a single associated device that sends data to the coordinator. The topology is symmetric; each sender is 3.5 m from both its own coordinator and the other PAN's coordinator (Fig. 2). In this stylized scenario, there are only two interfering transmissions: one where the interferer is at distance 3.5 m (the same distance as the sender) and one where the interferer is at distance 5 m. These distances are well within reliable communication range in our simulation, where the observed packet reception rate is nearly 100% for distances < 15 m (without interference). The compact topology also ensures that all senders reliably detect any ongoing transmission, i.e. there are no hidden terminals.

Having only one sender in each network simplifies interpretation of simulation results, since there is no in-network contention. Nevertheless, we wish to represent a realistically heavy load, such as might be generated by several senders that each generate one frame per beacon interval. The sender therefore generates traffic at a fixed rate of seven frames per beacon interval[1]. Traffic is periodic, but asynchronous with respect to the beacon, with a randomly chosen offset in each simulation run. Each frame is acknowledged by the coordinator and three transmission attempts are allowed per frame. The sender buffers up to 32 frames.

While this traffic model does not perfectly represent the timing behavior of frames generated by multiple senders, especially for slotted CSMA, it does capture key factors:

- Traffic is bursty, with multiple frames pending at the beginning of the active period. This occurs with random traffic patterns, where most frames are generated during

[1] We choose seven partly because the IEEE specification allows a maximum of seven senders to obtain GTS allocations for their transmissions.

OMNeT++ version	4.3.1
Castalia version	3.3
Tx power	-5 dBm
Path loss	α = 2.4, σ = 0, 55 dBm loss @ 1 m
min, max backoff	3, 5
Max backoffs	4
Max attempts	3
ACK's	yes
Short addresses	yes
Data frame size	121 + 6 bytes
MAC headers (MHR/MFR)	11 bytes
Aux. security header	10 bytes (4 byte MIC)
Data frame payload	100 bytes
CSMA traffic (7 frames)	48.64 ms (mean), 53.76 ms (max)
GTS traffic (7 frames)	38.08 ms
Buffer size	32 frames
ShortActive	
Beacon interval	983.04 ms (beacon order = 6)
Active period	61.44 ms (superframe order = 2)
Duty cycle	6.25 %
GTS allocation	46.08 ms (12 slots)
LongActive	
Beacon interval	983.04 ms (beacon order = 6)
Active period	122.88 ms (superframe order = 3)
Duty cycle	12.5 %
GTS allocation	84.48 ms (11 slots)

TABLE I: Experiment configuration. All IEEE 802.15.4 parameters are specified defaults.

the inactive (sleep) period, as well as with devices where a sensor is read just prior to turning on the radio.

- Intervals between slotted CSMA transmissions are irregular, with relatively long intervals between transmissions. Senders defer their transmissions if the channel is busy.
- GTS transmissions are highly regular, with transmissions closely packed at the end of the superframe, regardless of when frames are generated. Senders do not sense the channel or defer before transmitting.

B. Beacon interval and active period.

Interference between networks can only occur if the offset between their beacons is such that both networks are active at the same time. For this to occur, the beacon offset must be less than the duration of the active period, independent of the beacon interval separating active periods. Without loss of generality, we therefore set the beacon interval to 0.983 s, which gives a plausible duty cycle of 6-12% for our chosen traffic load.

For simplicity, we define two identically configured networks, each using either contention-based slotted CSMA or contention-free GTS to send data. In a slotted CSMA network, data transmissions can occupy the entire active period. In a GTS network, data transmissions can only take place during the GTS allocation; the active period must also include a contention period for broadcast and management traffic. Since our scenarios have only a single sender, there is a single, static GTS allocation.

We vary the length of the active period and the GTS allocation to define two configurations, which differ in their energy consumption and expected resilience to interference (Table I).

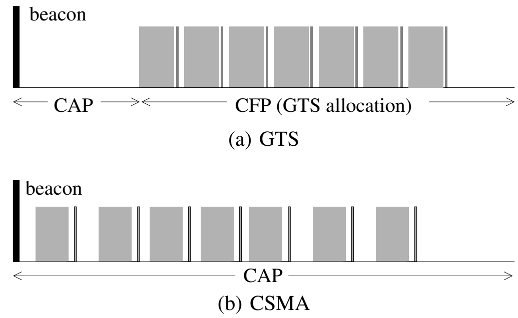

Fig. 3: ShortActive: The active period accommodates the beacon and seven data frames with acks, plus margin for one retransmission, even in the worst case CSMA backoff scenario.

ShortActive (Fig. 3): The active period and GTS allocation are defined such that they are just long enough to accommodate the network load. In a GTS network, one additional transmission is possible during the GTS allocation. Using slotted CSMA, at least one additional backoff and transmission attempt is possible (more in the case of favorable random backoffs). Because it allows for a small number of additional transmissions, ShortActive provides some resilience to occasional loss of a packet. It represents a conventional configuration for very low energy consumption networks.

LongActive: This configuration doubles the duration of the active period, as well as the GTS allocation. The GTS allocation is long enough to allow two transmission attempts for every frame. In a slotted CSMA network, there will be time for at least one additional backoff and re-transmission for every frame. In principle, LongActive is therefore able to accommodate the traffic load of both networks. Alternatively, LongActive allows for sending frames that were not able to be sent during previous beacon intervals and were buffered. However, this increased resilience comes at the cost of increased energy consumption, which is roughly proportional to the duration of the active period and also doubles.

C. Castalia

Castalia [5] is a widely used simulator for wireless sensor and body area networks and is written using OMNeT++ [6], an open source discrete event simulation environment. Castalia provides simulation models for the IEEE 802.15.4 PAN, including both slotted CSMA and GTS allocation[2].

To focus on inter-network interaction, we use a simple path loss model without fading to model the signal propagation. In general, we use Castalia's default frame reception model, where the probability of bit error depends on the received SINR and a frame is lost if any of the bits are in error. These behaviors are characterized in Fig. 2b. To better understand the extent to which our results depend on the reception model, some experiments were repeated using a simple SINR threshold reception model.

[2]We fixed several relatively minor issues in the Castalia code and expect to make changes available in collaboration with the Castalia developers.

The current Castalia implementation does have one limitation that is relevant for our results: It provides only an abbreviated model of the association process. In particular, it does not include the full exchange of frames for requesting and confirming association, assigning short local addresses, and revoking allocations following disassociation. This suggests that our results may somewhat overestimate the ease with which a node can re-associate with its PAN coordinator and resume sending application data after it becomes disassociated following the loss of several consecutive beacons.

IV. Results and Discussion

In Sections IV-A–IV-C, we present simulation results showing three interaction scenarios: between networks using contention-free GTS allocations, between networks using contention-based slotted CSMA, and between networks using both methods. We evaluate these interactions using two different duty cycles, the ShortActive and LongActive configurations.

Each interaction is characterized by the time offset between the beacon transmissions of the two coordinators and we use the packet reception rate as our performance metric. The beacon offset determines how much (if any) of the two networks' active periods overlap in time, as well as which parts of each active period (and hence which frames) are most vulnerable to interfering transmissions. For most offset values, one network's active period occurs during the other network's inactive period and there is no interaction effect. These offsets are not shown, but their impact is discussed in Section IV-D.

The simulation data are shown in Figs. 4–6. For each offset, the simulation is run for 1010 s (approximately 1000 beacon intervals and 7000 transmissions, plus a 10 s warmup period to ensure that association and GTS allocation are complete). Offsets are measured with a granularity of 492 μs, which is smaller than all possible gaps between transmissions (except the turnaround time between a data frame and ack). Each simulation was repeated 100 times. The dark line shows the average packet reception rate and the light shaded areas show the max and min over 100 repetitions.

For each scenario, we identify some of the key interference behaviors. We use the traffic representation in Fig. 3 to explain these behaviors and show that the loss of beacons due to collision is the most significant factor in overall packet loss. Comparing the performance of the ShortActive and LongActive configurations, we show that even doubling the active period (and energy consumption) does not necessarily lead to a large improvement in packet reception rate. However, when the longer active period does provide useful (re-)transmission opportunities, the ability to buffer packets plays an important role in recovering from beacon loss. Finally, in Section IV-D, we consider the implications of these results for the dynamic behavior of co-located networks. We describe three modes of interaction that are significant for network performance.

A. *CSMA-CSMA interaction*

Fig. 4 shows the interaction between two networks (CSMA-0 and CSMA-1), both using contention-based slotted CSMA. Since the two networks are the same, the interactions between them are symmetric with respect to the beacon offset and we only discuss $\text{offset} > 0$.

Figs. 4c and 4c show the case where the CSMA-0 beacon is received without interference, but the CSMA-1 beacon is sent while CSMA-0 frames are being transmitted. Since the beacon is sent without sensing the channel or deferring, it can be lost due to collision. If the beacon is lost, the CSMA-1 network cannot send in this beacon interval. The CSMA-0 packets experience no contention and the ShortActive configuration allows for re-transmission of the CSMA-0 frame that may have been lost due to collision with the CSMA-1 beacon. The CSMA-0 network has packet reception rate of 1 and the CSMA-1 network has packet reception rate 0.

However, there is a significant probability that the CSMA-1 beacon is successfully received. Not only are the CSMA-0 data transmissions proceeded by a random backoff, the CSMA-0 sender will also detect and defer to an ongoing CSMA-1 beacon transmission. The average backoff is 1.24 ms and the maximum is 2.24 ms (including the interframe space) and the contention window is 320 μs, leaving a window of opportunity for the CSMA-1 beacon (800 μs). In this case, networks CSMA-0 and CSMA-1 share the channel using (non-slotted) CSMA.

For any given beacon offset, the probability that the CSMA-1 beacon is successfully received is the probability that there is a gap in CSMA-0 transmissions at that time. If it is the i-th gap, this probability depends on the random backoffs of the $i-1$ previous frames. This probability distribution of the CSMA-1 beacon being interleaved with CSMA-0 traffic is reflected in Fig. 4a. As i grows (i.e. the offset increases) the probability of a combination of backoffs that lead to a gap at a particular offset also increases. Thus the peaks get slightly wider and the minimum packet reception rate increases as the offset increases.

Increasing the length of the active period does not change the underlying behavior. The driving factor is the probability that the CSMA-1 beacon is successfully received and this does not depend on the active period. However, in the LongActive configuration, both the peaks and minimums are higher. The longer active period and the use of CMSA lets both networks take advantage of the additional transmission opportunities. In particular, the CSMA-1 network can send packets that were buffered during beacon intervals in which the beacon was not successfully received.

The simulation results are very consistent, with little variation between the maximum and minimum packet reception rate over 100 simulation runs. The exception is in the latter part of the active period. Although the average packet reception rate is nearly 1, the minimum is 0. For the most part, this is an artifact of our traffic model. Traffic is periodic, but asynchronous (and random) with respect to the beacon offset.

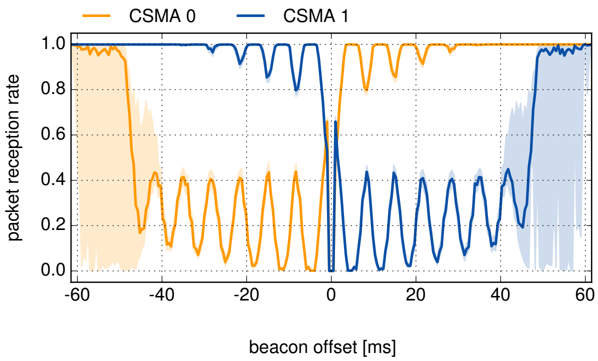

(a) CSMA-CSMA (ShortActive)

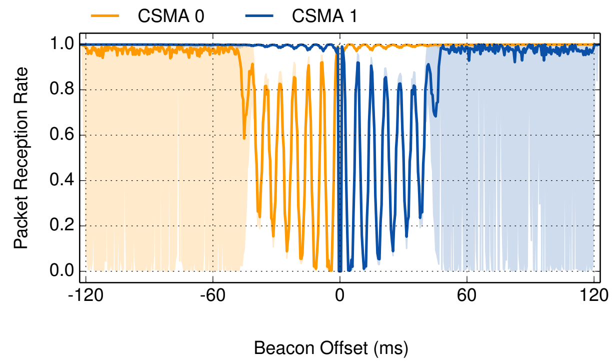

(b) CSMA-CSMA (LongActive)

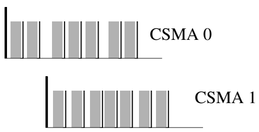

(c) CSMA-1 beacons interleaved with CSMA-0 backoff

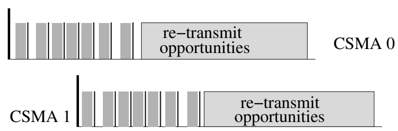

(d) CSMA-1 beacons interleaved with CSMA-0 backoff

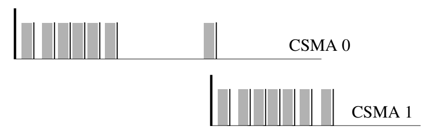

(e) packet arrival later in the active period

Fig. 4: CSMA-CSMA

Since the active period is short compared to the beacon interval (6.25 % and 12.5 % duty cycles), with high probability all seven packets arrive during the inactive period and are pending at the beginning of the active period. But it is also possible that a packet arrives during the active period, as in Fig. 4e.

If a CSMA-0 packet arrives late in the CSMA-0 active period, it will be transmitted immediately, using the usual sensing and backoff procedure. If the beacon offset is such that the CSMA-1 beacon collides with this frame (Fig. 4e), then the CSMA-1 network cannot send during the beacon interval and has 0 packet reception rate. Since the traffic has a random offset relative to the beacon in each simulation, this situation is likely to occur at least once over many simulation runs, leading to the observed variation. For any *specific* network instance, however, the packet reception rate will be 1, except for a narrow range of beacon offsets that are affected by the "stray" packet. A similar effect is observed in the GTS-GTS and CSMA-GTS scenarios as well.

B. *GTS-GTS interaction*

Fig. 5 shows the interaction between two networks (GTS-0 and GTS-1), both using contention-free GTS allocations. As in the CSMA-CSMA case, the interaction is symmetric and we only consider $\mathrm{offset} > 0$.

Figs 5c and 5d show the cases where both network's beacons are received without interference, but the two networks' subsequent GTS allocations overlap. The GTS senders cannot defer to each other and their transmissions occupy the channel almost continuously, with only a fixed interframe space (IFS) of 640 μs. As a result, frames in both networks are lost due to collision. Only frames that are transmitted at the beginning (in GTS-0) or end (in GTS-1) of each network's GTS allocation avoid interference and are successfully received.

This results in the "step function" starting at $\mathrm{offset} = 0$: In the ShortActive configuration (Fig. 5a), first one and then two frames in each GTS allocation are received without interference as the beacon offset increases. In the LongActive configuration (Fig. 5b), there is a larger range of offsets over which this situation occurs and thus a larger number of "steps". However, the much larger number of transmission opportunities in the LongActive configuration itself does *not* improve performance – both networks simply spend the time uselessly transmitting frames that are doomed to be lost.

At larger beacon offsets (Figs. 5e and 5f), the GTS-1 beacon is sent during the GTS allocation of network GTS-0. Since neither beacons nor GTS frames defer to ongoing transmissions, the GTS-1 beacon is (usually) lost due collision with GTS-0. If this happens, GTS-1 cannot send at all during the beacon interval. This allows GTS-0 to send its frames without interference. In the ShortActive configuration, the GTS allocation allows for one re-transmission, so the GTS-0 frame that collided with the GTS-1 beacon can be re-transmitted, if needed. As a result, there is almost no packet loss in GTS-0 and almost complete packet loss in GTS-1.

If the GTS-1 beacon is transmitted in the short IFS between GTS-0 transmissions, there is a very small probability that this very short frame is successfully received despite interference. Both senders then attempt to use their GTS allocations, with limited success, as in the case described above in Fig. 5c. This results in the small matched peaks and dips; there is one for each gap between the GTS-0 frames. In the LongActive scenario, the peaks are much larger and the dips smaller than

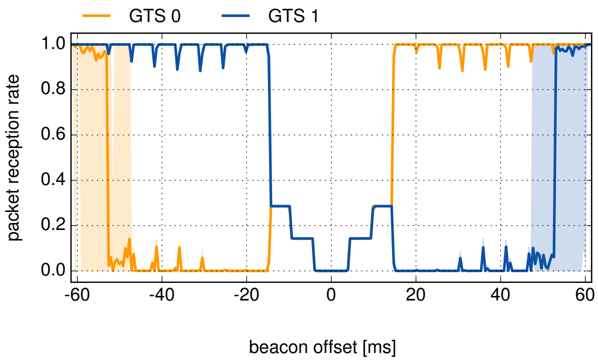

(a) GTS-GTS (ShortActive)

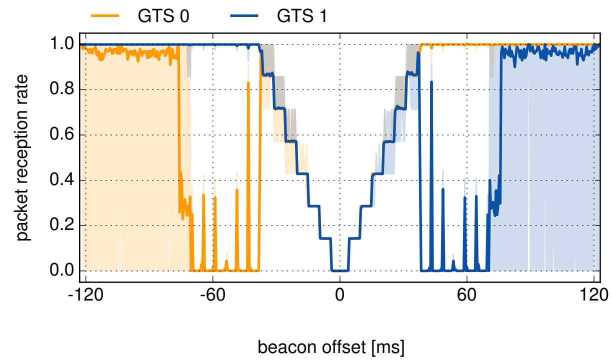

(b) GTS-GTS (LongActive): Note the different x-axis.

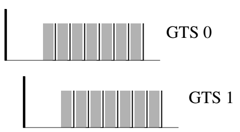

(c) packet loss

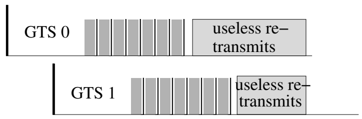

(d) packet loss and useless re-transmits

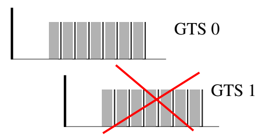

(e) GTS-1 beacon often lost

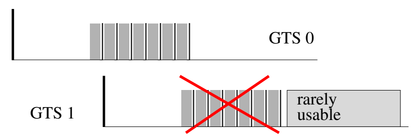

(f) GTS-1 beacon often lost

Fig. 5: GTS-GTS

in the ShortActive scenario. Although the GTS-1 beacon is no more likely to be successfully received in the LongActive configuration, the longer GTS allocation gives the G TS-1 sender additional transmit opportunities. Not only can it transmit frames from this beacon interval, it may also transmit frames that were buffered during previous beacon intervals in which it did not receive the beacon.

This effect is highly sensitive to the simulation's probabilistic frame reception model, which determines the probability that a frame is received despite interference and does not appear in simulations based on a simple threshold frame reception model. On the other hand, this effect may be even more likely to occur in real wireless environments, with their often variable link quality.

C. CSMA-GTS interaction

Fig. 6 shows the interaction between two networks, one using contention-based unslotted CSMA and the other using contention-free GTS allocations.

Figs 6c and 6d show the case of positive offsets. It is very similar to the CSMA-CSMA case in Fig. 4, where the behavior reflects the probability distribution of the CSMA-1 beacon being interleaved with CSMA-0 traffic. Here, it is the GTS beacon that is interleaved with CSMA traffic. If the GTS beacon is lost, the GTS sender cannot transmit during that beacon interval. The CSMA packet reception rate is 1 and the GTS packet reception rate is 0.

If the GTS beacon is successful, the CSMA and GTS networks share the channel. Because the GTS network cannot defer, the two networks share the channel somewhat less effectively than two CSMA networks do. This is visible in the LongActive configuration: In the CSMA-CSMA scenario, CSMA-1 obtains slightly higher peak and higher minimum packet reception rates when interacting with CSMA-0 (Fig. 4b) than the GTS network obtains when interacting with the CSMA network (Fig. 6b).

Fig. 6e shows the case of small negative offsets. Both beacons are always successful and the CSMA and GTS networks share the channel in every beacon interval. In the ShortActive configuration, the CSMA sender must defer to the GTS sender, which obtains a high packet reception rate. The CSMA sender is generally able to send only two frames, either before or after the GTS transmissions. The LongActive configuration provides enough (re-)transmission opportunities for both networks to achieve high packet reception rates, although the CSMA network is still affected by the more aggressive GTS network (Figs 6b and 6e).

Fig. 6f shows the case of larger negative offsets. The CSMA beacon is usually lost due to collisions with GTS transmissions and no CSMA data is sent in those beacon intervals. There is a very small probability that the CSMA beacon is received despite interference and the two networks share the channel during that beacon interval. This results in the small matching peaks and dips, similar to those in Figs. 5a and 5b for the GTS-GTS case. The LongActive configuration allows the CSMA networks to take advantage of its additional transmission opportunities and obtain higher packet reception rates.

D. Discussion and implications for dynamic behavior

The simulation results show that inter-network interaction leads to severe performance degradation in all three interaction scenarios. This is true even in cases where each network's active period is, in principle, large enough to accommodate all of the traffic in both networks.

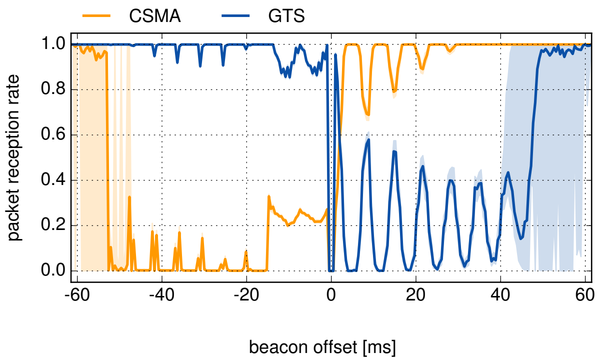

(a) CSMA-GTS (ShortActive)

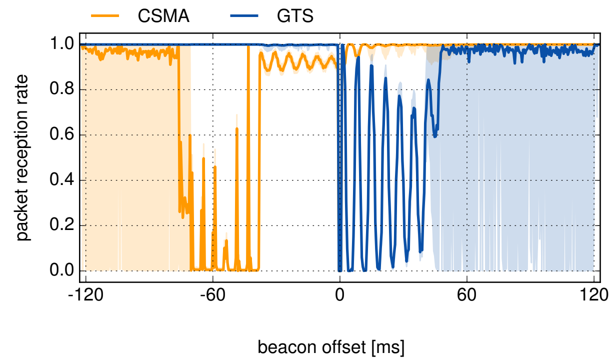

(b) CSMA-GTS (LongActive)

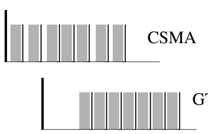

(c) GTS beacons interleaved with CSMA backoff

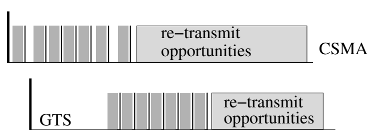

(d) GTS beacons interleaved with CSMA backoff

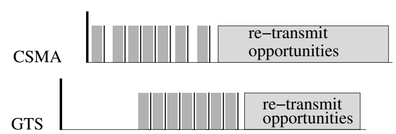

(e) CSMA and GTS traffic share channel

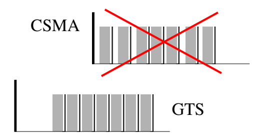

(f) CSMA beacons lost

Fig. 6: CSMA-GTS

Although the scenarios differ in specific behaviors, there are a number of common elements. In particular, the loss of beacon frames has a very large impact on the overall packet reception rate. Most importantly, the probability of beacon loss is highly dependent on the offset. In cases where there is a non-trivial probability of beacon reception, the additional transmission opportunities provided by a longer active period can provide some improvement in the packet reception rate. However, these improvements are generally moderate compared to the energy cost of doubling the length of the active period.

The results also have implications for the networks' dynamic behavior. The key insight is that co-located networks will experience *all* possible beacon offsets over time due to the clock drift between them. In beacon-enabled PANs, associated devices are tightly synchronized with their coordinator via the beacon. But the clocks belonging to coordinators in independent PANs will drift relative to one another. The rate of drift will depend on a number of factors: The IEEE 802.15.4 specification requires the radio to have an underlying accuracy of at least $\pm$ 40 ppm. At the other extreme, two PAN coordinators might both be synchronized with an external UTC time source to within a couple of ppm.

Figs. 4 - 6 can therefore be interpreted as showing the behavior of the two networks as their beacon offset changes from -61.44 ms to +61.44 ms (or -122.88 ms to +122.88 ms) due to clock drift. Assuming a relative clock drift of 10 ppm, the inter-network interaction would last about 3.5 hours, after which the networks' active periods would not overlap again until approximately 24 hours later. We can therefore translate our results into a description of network dynamics, identifying three fundamental modes of interaction.

The first is oscillation, which would be experienced in the cases shown in Figs. 4a, 4b, 6a, and 6b). The offset-dependent probability that one network's beacons are successfully interleaved with another network's CSMA transmissions will result in large, slow oscillations in the packet reception rate over time. The amplitude of the oscillations depends in part on the length of the active period. Assuming a 10 ppm relative clock drift, networks could experience episodes of over an hour during which the packet reception rate would cycle from very high to very low and back again with a period of some 10 min.

The second mode of interaction is slow deterioration and recovery, which would be experienced in the GTS-GTS scenarios, especially Fig. 5b, where the network will experience a stepwise deterioration in packet reception rate over roughly an hour, followed by some 16 min of complete packet loss, and a similar stepwise recovery.

The third mode of interaction is bursty. In some cases, the beacon is received only rarely. Particularly with longer active periods, the sender can take significant advantage of these opportunities, as shown in Figs. 5b and 6b. This would be experienced as episodes of an hour or more with almost complete packet loss, interspersed with bursts of recovery lasting some 10 s out of every 10 min.

Conversely, periodic traffic patterns can result in the presence of "stray" packets, such as shown in Fig. 4e. Such situations would be experienced as an episode of up to several minutes of partial or even complete beacon and/or packet loss, during extended periods of otherwise uninterrupted operation.

These dynamic behaviors also have implications for the senders ability to buffer traffic. The total packet reception rate over time will depend not only on the transmission opportunities in the active period, but also the sender's capacity to buffer packets, relative to the frequency of transmission opportunities.

V. Related Work

Many works address interferer classification and interference mitigation, especially cross-technology interference such as between IEEE 802.11 and IEEE 802.15.4. Practical interferer identification based on learning methods applied to characteristic error patterns is reported in [7]. More generally, multi-channel operation allows networks to avoid interference by finding the best available channel or by hopping between channels. IEEE 802.15.4e [8]/TSCH (time slotted channel hopping) allows PANs to coordinate their channel selection, but does not specify a channel hopping policy. These approaches are effective only as long as there exist channels that have low interference. More complex interference resilient protocol design approaches include [9] and [10].

Early work on IEEE 802.15.4 interactions [11] considered the problem of "semantic interference", which occurs when one network receives packets that originate in another network. That work emphasizes the need to authenticate all frames to prevent "foreign" frames from being accidentally misinterpreted, but it also speculates about the possibility of the kinds of protocol-level interactions studied in this work.

Perhaps the only other direct study of IEEE 802.15.4 interactions is [12]. In that work, the beacon offset is also used to characterize the interaction scenarios, which focus on whether it is senders or coordinators that experience interference. The authors also used mobility to create complex scenarios with variable interference. As in our work, the authors observe significant losses and highlight beacon loss as the key factor. However [12] was based on very coarse grain differences in beacon offset. By contrast, our measurements have more than two orders of magnitude finer granularity ($< 500\mu s$) and identify significantly more complex behaviors, particularly dynamic behaviors in the presence of clock drift. We also study the tradeoff between energy consumption and performance.

Other relevant work is found in the IEEE 802.15.6 [13] community, where reliability and interference mitigation are important issues for body area networks. The IEEE 802.15.6 standard has proposed "beacon shifting" to mitigate the problem of persistent beacon loss. This allows the beacon to be transmitted at times other than at the beginning of the beacon interval, to increase resilience to beacon loss [14].

VI. Conclusion

We present a simulation study of inter-network interaction between two IEEE 802.15.4 beacon-enabled PANs using contention-based slotted CSMA and contention-free GTS allocations for communication. Using stylized scenarios to highlight the fundamental modes of interaction between network, we identify and explain a number of behaviors that may have significant implications for protocol and application design.

These include the possibility of episodes of large, slow oscillations in packet reception rate, slow stepwise deterioration and recovery, brief bursts of moderately high packet reception rate in the midst of extended periods of almost complete packet loss, and short periods of high packet loss during periods of otherwise uninterrupted operation. Some of the factors driving these behaviors include the time- (or offset)-dependent probability of beacon loss, the amount and quality of (re-)transmission opportunities, and the sender's ability to buffer packets. These factors also combine to limit the benefit that can be obtained through even significant increases in the active period and hence the energy consumption.

In real environments, the effects we discuss would be mingled with many others, including in-network contention and interference, more variable traffic patterns, and changes in the number and kinds of networks present. There would also be more variability in link quality, with each device in the network experiencing different levels of interference from various senders.

Even though our results consider only a very simple network scenario, they already show that these interactions can be extremely complex. The design of protocols and application that are intended to work in realistic IoT environments therefore needs to be supported by rigorous, detailed performance evaluation that takes into account realistic models of inter-network interaction.

References

[1] A. Dunkels, "The ContikiMAC Radio Duty Cycling Protocol," Swedish Institute of Computer Science, Tech. Rep. T2011:13, December 2011.

[2] The ZigBee Alliance, "http://www.zigbee.org."

[3] Z. Shelby and C. Bormann, *6LoWPAN: The Wireless Embedded Internet*. John Wiley & Sons, 2011.

[4] HART Communication Foundation, "http://www.hartcomm.org."

[5] A. Boulis, "Castalia: Revealing Pitfalls in Designing Distributed Algorithms in WSN," in *5th Int'l Conference on Embedded Networked Sensor Systems (SenSys'07)*, 2007.

[6] A. Varga, "OMNeT++," in *Modeling and Tools for Network Simulation*, K. Wehrle, M. Günes, and J. Gross, Eds. Springer-Verlag, 2010.

[7] V. Iyer, F. Hermans, and T. Voigt, "Detecting and avoiding multiple sources of interference in the 2.4 GHz spectrum," in *12th European Conf on Wireless Sensor Networks (EWSN)*, 2015.

[8] "Part 15.4: Low-rate wireless personal area networks (LR-WPANs) amendment 1: MAC sublayer," *IEEE Std 802.15.4e-2012*, 2012.

[9] I. Glaropoulos, V. Fodor, L. Pescosolido, and C. Petrioli, "Cognitive WSN transmission control for energy efficiency under WLAN coexistence," in *6th Int'l Conf on Cognitive Radio Oriented Wireless Networks and Communications*, 2011.

[10] Q. Javed and R. Prakash, "Chameleon: A framework for coexistence of wireless technologies in an unlicensed band," *Wireless Personal Communication*, vol. 77, no. 1, 2014.

[11] L. M. Feeney, "Exploring semantic interference in heterogenous sensor networks," in *1st ACM Workshop on Heterogeneous Sensor and Actor Networks*, 2008.

[12] N. Nordin and F. Dressler, "Effects and Implications of Beacon Collisions in Co-located IEEE 802.15.4 networks," in *IEEE Vehicular Technology Conference (VTC Fall)*, 2012.

[13] "IEEE Standard for Local and metropolitan area networks - Part 15.6: Wireless Body Area Networks," *IEEE Std 802.15.6-2012*, 2012.

[14] S. Kim, S. Kim, J.-W. Kim, and D.-S. Eom, "A beacon interval shifting scheme for interference mitigation in body area networks," *Sensors*, vol. 12, no. 8, 2012.