

The Impact of the Physical Layer on the Performance of Concurrent Transmissions

Michael Baddeley*, Carlo Alberto Boano[†], Antonio Escobar-Molero[‡], Ye Liu[§], Xiaoyuan Ma[¶],
Usman Raza*, Kay Römer[†], Markus Schuß[†], and Aleksandar Stanoev*

*Toshiba Europe Ltd., Bristol, United Kingdom – {michael.baddeley,usman.raza,aleksandar.stanoev}@toshiba-bril.com

[†]Institute of Technical Informatics, Graz University of Technology, Austria – {cboano,markus.schuss,roemer}@tugraz.at

[‡]RedNodeLabs UG, Munich, Germany – antonio@rednodelabs.com

[§]College of Engineering, Nanjing Agricultural University, China – yeliu@njau.edu.cn

[¶]Shanghai Advanced Research Institute, Chinese Academy of Sciences, China – maxy@sari.ac.cn

Abstract—The popularity of concurrent transmissions (CT) has soared after recent studies have shown their feasibility on the four physical layers specified by BLE 5, hence providing an alternative to the use of IEEE 802.15.4 for the design of reliable and efficient low-power wireless protocols. However, to date, the extent to which physical layer properties affect the performance of CT has not yet been investigated in detail. This paper fills this gap and provides the first extensive study on the impact of the physical layer on CT-based solutions using IEEE 802.15.4 and BLE 5. We first highlight through simulation how the impact of errors induced by de-synchronization and beating on the performance of CT highly depends on the choice of the underlying physical layer. We then confirm these observations experimentally on real hardware through an analysis of the bit error distribution across received packets, unveiling possible techniques to effectively handle these errors. We further study the performance of CT-based flooding protocols in the presence of radio interference on a large-scale, and derive important insights on how the used physical layer affects their dependability.

I. INTRODUCTION

A recent breakthrough in the low-power wireless community has been the development of communication protocols based on Concurrent Transmissions (CT), which intentionally let multiple relaying nodes forward packets by simultaneously broadcasting them on the same carrier frequency. Thanks to the capture effect [1] and to non-destructive interference [2], nodes overhearing these concurrent transmissions have a high probability to receive at least one transmission correctly, enabling the creation of reliable and efficient cyber-physical systems and Internet of Things (IoT) applications [3].

The key benefit of CT is the ability to exploit sender diversity to realize simple flooding and synchronization services across large-scale multi-hop wireless networks [4]. Relaying nodes in a mesh network utilizing CT-based protocols do not need to explicitly avoid collisions using conventional techniques such as carrier sensing, and can avoid the overhead of routing and link-based communication [3].

A large body of work has proposed CT-based data collection [5], [6], [7] and dissemination [8], [9] protocols that can achieve unprecedented gains in terms of reliability, end-to-end latency, and energy efficiency. These protocols can outperform existing solutions even in the presence of harsh radio interference [10], [11], [12], [13], as shown by four editions of the EWSN dependability competition [14].

However, the vast majority of these protocols have only been implemented and verified experimentally using off-the-shelf platforms based on 2.4 GHz IEEE 802.15.4 radios (e.g., the popular but rather outdated TelosB mote [15]). These solutions employ a physical layer (PHY) based on Orthogonal Quadrature Phase Shift Keying (OQPSK) and Direct Sequence Spread Spectrum (DSSS), as specified by the IEEE 802.15.4 standard [16], where DSSS provides the coding robustness needed by CT to be sufficiently reliable [17].

CT and the impact of different PHYs. An experimental study by Al Nahas et al. [18] has shown the feasibility of CT also when using Bluetooth Low Energy (BLE). Their preliminary results show that reliable and efficient CT-based flooding is possible on BLE-based mesh networks, but highlight that the performance *largely depends on the employed PHY*, as confirmed by the measurements reported in [19], [20]. Indeed, the most recent version of Bluetooth Low Energy (BLE 5) supports four PHYs that largely differ in terms of data rate and robustness [21]: 2M (2Mbps), which doubles the nominal throughput of the original 1M PHY (1Mbps), and two coded PHYs with coding rates of 1/2 and 1/8 respectively (i.e., the 500K and 125K PHYs).

These preliminary observations are important in light of the increasing number of commodity IoT platforms that embed low-power radios supporting multiple wireless standards; hinting that developers need to carefully select the physical layer used for CT-based communication. Examples of such single-chip, multi-PHY platforms are the TI CC2652R [22] and the Nordic nRF52840 [23], which support 2.4 GHz IEEE 802.15.4 OQPSK-DSSS alongside the four BLE 5 PHYs.

However, to date, *the extent to which physical layer properties affect CT-based solutions employing IEEE 802.15.4 and BLE 5 has not yet been investigated in detail*. Firstly, there is no experimental study systematically analyzing how physical layer effects such as beating (a pulsating interference pattern between two or more signals at slightly different frequencies) induced by relative carrier frequency offset [2], and de-synchronization due to clock drift [24] affect the reliability of the received signal in the presence of multiple concurrent transmitters. Furthermore, there is no experimental work studying whether the performance of CT in harsh RF environments differs depending on the underlying PHY.




Fig. 1: Beating due to relative oscillator frequency inaccuracies between devices. Signals combine to produce periods of *constructive* and *destructive* interference (CI and DI).

Shedding light on these aspects is important to (i) provide a better understanding on the role of the physical layer on the reliability and efficiency of CT, as well as to (ii) empower developers to use the physical layer as a means to fine-tune the performance of CT-based protocols at runtime.

Our contributions. This paper represents the first in-depth experimental study on the impact of the PHY on CT-based solutions employing IEEE 802.15.4 and BLE 5, providing valuable observations intended to inform and direct the engineering of future CT-based network protocols.

Firstly, there exists no other work extensively examining and empirically proving the significance of the beating effect on CT performance, while the highly novel technique we introduce to achieve this can potentially allow observation in real-time, without a Software Defined Radio (SDR). Secondly, while recent selected works have demonstrated the feasibility of CT over BLE in single-hop experiments [18], [19], [20], [25], this paper is the first to systematically examine the impact of different PHYs in a multi-hop scenario, and under benchmarkable interference. With modern chips introducing real-time PHY switching capabilities with *no additional radio overhead* [22], [23], these insights can provide a blueprint for the design of CT-based networking protocols capable of switching PHY at runtime in response to changing network conditions towards a more dependable performance.

We start by simulating the performance of CT for the different BLE 5 PHYs, highlighting the role of beating under different interference scenarios. We then set up an extensive experimental campaign to confirm these results and to systematically study CT performance across all IEEE 802.15.4 and BLE 5 PHYs supported by the nRF52840 platform. To this end, we use the D-Cube public testbed [26], [27], recently enhanced with 50 nRF52840-DK devices, to observe both beating frequencies and de-synchronization effects on real hardware through analysis of the error distribution across received packets. Our experiments demonstrate that the impact of errors induced by de-synchronization and beating on CT performance is highly dependent on the choice of the underlying PHY, on the relative carrier frequency offset between transmitting devices, as well as on the number of concurrent transmitters. Specifically, we observe that: (i) high data rate PHYs experience wider beating and we can mitigate its impact through repetition; (ii) if the power delta between signals is insufficient, then BLE 5 convolutional coding is does not sustain reliable CT; (iii) the pattern mapper used in BLE 5 125K allows it to effectively handle narrow beating.




Fig. 2: Beating manifests in many forms depending on the relative frequencies between devices and their channel gains. Strong beating attenuates a signal toward zero in a beating ‘valley’.

We further use D-Cube to perform the first experimental study on the performance of different CT-based flooding protocols as a function of the underlying PHY in the presence of RF interference on a large scale. To this end, we make use of D-Cube’s JamLab-NG functionality [28] to generate artificial Wi-Fi interference and stress-test the performance of CT-based protocols such as Glossy [4] and robust flooding (RoF) [10] under *no*, *mild*, and *strong* interference. Our results allow us to derive important insights on which PHYs are effective to help CT-based protocols in mitigating the impact of interference. Such insights include: (i) the superiority of IEEE 802.15.4 and BLE 5 500K PHY under strong interference, (ii) the fact that the BLE 125K PHY should not be used in conjunction with long payload lengths under interference, as well as (iii) the need to dynamically change PHY at runtime to provide the best trade-off between reliability, latency, and energy efficiency.

After providing some background knowledge on CT in Sect. II, this paper makes the following specific contributions:

- We simulate the performance of CT for all BLE 5 PHYs, highlighting the role of beating (Sect. III).
- We are the first to experimentally observe beating frequencies and de-synchronization effects on real hardware and wireless channel for different PHYs through an analysis of the bit error distribution across received packets (Sect. IV).
- We evaluate the performance of CT-based flooding protocols under radio interference, and provide insights on how the employed PHYs affect dependability (Sect. V).

We then describe related work in Sect. VI to highlight how our insights align with existing literature, and conclude our paper in Sect. VII along with a discussion on future work.

II. CONCURRENT TRANSMISSIONS IN LOW-POWER WIRELESS NETWORKS

CT is the concept by which several nodes transmit the data they want to share at the same time. CT-protocols have typically been based on different variations of binary frequency-shift keying (BFSK)¹, as specified in the BLE 5 and IEEE 802.15.4 standards (for a comprehensive overview, please refer to [16], [21], [30]). Early works such as Glossy [4] showed that, when using frequency-based modulations, if nodes are sufficiently synchronized, then transmissions of the *same data* will align and the packet will be correctly received with cooperative gain. Meanwhile, later works have shown that transmissions of

¹OQPSK with half-sine pulse shaping is equivalent to Minimum-Shift Keying (MSK) and can be demodulated as a frequency modulation [29].

different data greatly benefit from capture effect due to energy diversity between transmitters. CT hence constitute a robust technique to deploy simple and latency-optimal mesh networks. Nevertheless, recent literature has shown that CT introduce two types of errors that degrade communication performance:

- 1) *Synchronization errors.* Concurrent transmitters are not perfectly synchronized, which introduces intersymbol interference when different bits overlap on the air. To minimize this effect, packet transmissions must be triggered within a time interval ideally lower than half the symbol period [24].
- 2) *Beating Effect.* When CT overlap on the air, the resulting waveform has a beating amplitude due to alternating periods of constructive and destructive interference (beating). As shown in Fig. 1, with *two* concurrent transmitters, the waveform's envelope has a sinusoidal shape, while featuring more complex forms when *more than two* transmissions overlap [25]. While potentially introducing a certain degree of energy gain during peaks, beating greatly increases the chance of bit errors during low-energy periods (valleys) and has generally a negative net effect.

Notably, beating will impact dense topologies when there is no dominant transmission and will consequently be affected by deep fading. On the other hand, its impact is reduced when different transmissions are received with enough energy diversity and the capture effect kicks in. In Fig. 2 we categorize beating as *wide and strong*, *wide and weak*, *narrow and strong*, or *narrow and weak*. Beating will be randomly *narrow* or *wide* depending on the relative carrier frequency offset between transmitters, while it will manifest as *strong* or *weak* depending on the relative difference in received signal energy.

III. IMPACT OF PHYSICAL LAYER ON CT PERFORMANCE: ANALYSIS AND SIMULATION

While synchronization errors can be greatly reduced by properly designing the CT network protocol and its retransmission strategy, *beating cannot be avoided*. Beating always appears when signals from non-coherent transmitters overlap in the air, due to their different carrier frequency offset (CFO). Moreover, the temporal period of the beating is different for each set of concurrent transmitters and randomly depends on their oscillator inaccuracies. The unpredictable temporal length of the beating largely affects the error rate, since the beating periods can be very narrow (and several peaks and valleys can appear within a packet transmission), or very wide (and a packet can be completely shadowed within a valley).

To better analyze the impact that different PHYs have on the performance of CT under beating, we simulate the different communication systems using MATLAB to obtain the average Packet Error Rate (PER) vs. Signal-to-Noise Ratio (SNR) for two CT recovered with a non-coherent BFSK receiver, as in [31]. We assume constant additive white Gaussian noise and no synchronization errors. We repeat this for the different BLE 5 coded (500K and 125K) and uncoded (1M and 2M) PHYs, for different oscillator inaccuracies (which result in either wide or narrow beating) and power deltas. Both coded PHYs are based on the 1M PHY, adding a convolutional code of rate 1/2

and are received with a hard-decision Viterbi decoder [32]. In addition, the 125K PHY adds a Manchester pattern mapper of four elements per coded bit².

We define the Relative Frequency Offset (RFO) – which determines the beating frequency – as the difference between the CFO of each individual transmitter, and the power delta ΔP as the power ratio with which both CT are received. The SNR is defined relative to the strongest transmission (assuming $P_{R1} > P_{R2}$) and N being the noise power:

$$RFO(Hz) = |CFO_1 - CFO_2| = 1/T_{Beating}, \quad (1)$$

$$\Delta P(dB) = 10 \log_{10}(P_{R1}/P_{R2}), \quad (2)$$

$$SNR(dB) = 10 \log_{10}(P_{R1}/N). \quad (3)$$


The BLE standard requires the CFO to be within ± 150 kHz [30], which results in RFOs lower than 300 kHz. Therefore, the RFO is always lower than the (coded or uncoded) symbol frequency (i.e., 2 MHz in BLE 5 2M and 1 MHz for the other three PHYs).

The results of our simulation are presented in Fig. 3. Based on these results, we derive the following observations:


- 1) **Impact of beating.** We first compare the results obtained with two concurrent transmitters (2 CT) with those obtained with a single transmitter (no beating). With low-noise ($SNR > 15$ dB), beating negatively affects packet reception and increases the PER. Only when operating in high-noise conditions ($SNR < 10$ dB), 2 CT experience a PER lower than that of a single transmitter, due to the positive net effect of constructive interference intervals. CT are hence an optimal mechanism in harsh environments with high noise, in which packet loss is high. Otherwise, the effect of destructive interference dominates and the PER increases.
- 2) **Wide ($T_{Beating} > T_{Packet}$) and strong ($\Delta P \approx 0$ dB) beating, Fig. 3a.** In this case, T_{Packet} , which denotes the over-the-air time of a packet³, is the key factor dictating the PER. Indeed, the probability that the transmission spans a destructive interference interval is lower as the time the packet spends on the air decreases. Hence, when subjected to a same fixed $T_{Beating}$, uncoded PHYs (BLE 5 1M and 2M) perform better than coded ones (125K and 500K), since the former benefit from faster transmissions. With wide energy valleys, convolutional codes are ineffective: as a result, BLE 5 125K is the worst performing PHY, since it features the longest packet durations.
- 3) **Narrow ($T_{Beating} < T_{Packet}$) and strong ($\Delta P \approx 0$ dB) beating, Fig. 3c.** In this configuration, the packet transmission always spans one or more destructive valleys. The BLE 5 2M PHY benefits from having the shortest packet duration, outperforming the 1M PHY. The convolutional encoder used in the BLE 5 500K and 125K PHYs is ineffective against beating, since it is optimal for discretely distributed one bit errors, but not to correct the burst errors that typically occur with beating.

²When using BLE 5 125K PHY's Manchester Pattern Mapper, a '0' is translated into '0011', whereas a '1' is translated to '1100' [30], [33].


³ T_{Packet} is computed as $T_{Packet} = B \cdot (1/DR)$, with DR being the data rate of the chosen PHY and B being the length of the packet in bits.




(a) Wide and strong beating (2 CT, RFO=500Hz, $\Delta P=0$ dB)



(b) Wide and weak beating (2 CT, RFO=500Hz, $\Delta P=1$ dB)



(c) Narrow and strong beating (2 CT, RFO=10kHz, $\Delta P=0$ dB)



(d) Narrow and weak beating (2 CT, RFO=10kHz, $\Delta P=1$ dB)

Fig. 3: Simulation results showing the PER when sending 30-byte packets with two concurrent transmitters and with a single transmitter (no beating). Differences in beating periods (wide or narrow) and power deltas have a significant impact on how beating affects performance, which explains why different node pairs (with different RFOs) may experience very different PERs.

Nevertheless, the BLE5 125K PHY features a good performance in narrow beating conditions when noise is very low ($\text{SNR} > 20$ dB), experiencing a waterfall-like PER decrease, whereas BLE5 500K does not experience such a decrease until $\text{SNR} > 30$ dB, performing worse than uncoded PHYs in the mid-to-low noise range.

- 4) **Very low noise ($\text{SNR} > 25$ dB) and strong beating ($\Delta P \approx 0$ dB), Fig. 3a and Fig. 3c.** Under these conditions, the most promising PHYs are BLE5 2M, when fast data rate is needed, and BLE5 125K when long-range is required and a limited data rate is sufficient. With the BLE5 2M PHY, in a flooding-based mesh network using CT, the optimal strategy is using several (fast) retransmissions to increase the chances of successful packet receptions and compensate the lower PER compared to classical (without beating) routing schemes. In dense networks, in which multiple (more than two) retransmitters with different RFOs are simultaneously in the range of the receivers, narrow beating conditions are more frequent, and therefore one can exploit the performance boost of the BLE5 125K PHY, which potentially requires no retransmissions (in contrast to the use of BLE5 2M PHY). However, if packets are too long, BLE5 125K may enter the wide-beating region, in which it performs poorly, and may suffer from synchronization errors. Therefore, with CT, BLE5 125K should not be used to transmit packets with a large payload.
- 5) **Very high noise ($\text{SNR} < 5$ dB), Fig. 3a and Fig. 3c.** In such configuration, the BLE5 500K PHY performs well and represents the best choice to survive extremely noisy environments. Nonetheless, we believe that as the number of transmitters increases the properties of the system tend

to be dominated by the internal interference caused by beating, decreasing the influence of external noise.

- 6) **Weak beating ($\Delta P > 0$ dB), Fig. 3b and 3d.** In real deployments, CT are normally received with dissimilar power levels. Under weak beating, the signal does not completely fade during the valleys, which greatly decreases the impact of beating on the PER. For greater dissimilarities ($\Delta P > 3$ dB), the capture effect kicks in, and the PER becomes very similar to the no-beating scenario. Under weak beating, the 1M PHY may also be a suitable alternative, especially in sparse networks where beating is wider. This is because it has a comparable performance to the 2M PHY, while being more tolerant to synchronization errors (it requires a synchronization within $0.5 \mu\text{s}$ instead of $0.25 \mu\text{s}$), and having a better receiver sensitivity.

In real-world networks, all four scenarios depicted in Fig. 3 may simultaneously appear in different sections of a multi-hop network, depending on the practically unpredictable CFOs and power level relationships between the concurrent transmitters. Even for a given link, conditions may change over time, since temperature alters the CFO and surrounding interference or multipath propagation cause dynamic fluctuations in the power deltas. It is hence desirable that an optimal PHY tailored to handle CT features robust behavior in all four scenarios, since controlling operating conditions within a wireless network is challenging. In particular, within dense mesh networks with potentially many more than two CT at every hop, narrow and strength-varying beating is expected to dominate, and links can be modelled as behaving like the narrow beating figures for 2 CT (Fig. 3c and Fig. 3d) working in the high-SNR region.

Observations for beating mitigation. Based on these simulation results, we infer that beating can be mitigated by boosting energy diversity with proper node placement and techniques to dynamically control the transmission power. Nonetheless, this would affect the main advantages of using CT-based protocols: scalability and simplicity. When using lower data rates, which ultimately results in narrower beating periods relative to the packet period, it is more likely that packets will not completely fade: here, coding and Forward Error Correction (FEC) techniques can be exploited to introduce sufficient diversity to recover the errors introduced during the valleys by using the information received during the peaks. This is the case for IEEE 802.15.4, which uses Direct-Sequence Spread Spectrum (DSSS), and for the BLE coded PHYs (500K and 125K), which feature FEC. However, the convolutional encoder used in BLE coded PHYs is not effective in low-noise conditions, and the main gain in this region comes from the addition of the Manchester pattern mapper in BLE 5 125K. An exception is the case of networks operating in very noisy environments, in which the coded PHYs (especially BLE 5 500K) may be able to trigger successful packet receptions where classical (routing-based) schemes and coded PHYs may fail. Contrarily, higher data rates experience relatively wider beating periods, since the packet duration is shorter. In this scenario, whole packets may be blocked by wide valleys, which completely precludes any correction attempt. Similarly, the packet may randomly experience a wide peak, and be properly received. With high data rates, as in BLE uncoded PHYs (1M and 2M), the strategy should hence be different. Instead of trying to correct the errors, it would be more effective to repeat the packet several times, practically trying to randomly trigger a transmission spanning an energy peak.

Next, we confirm our simulation results on real hardware using *over-the-air* experiments. As we expect physical layer properties to cause a fluctuating error probability across received packets, we map the distribution of bit errors across a packet so that, practically, we observe the beating frequency over time according to the data rate of the underlying PHY.

IV. CT PERFORMANCE OVER DIFFERENT PHYs: EXPERIMENTAL EVALUATION OF BIT ERROR DISTRIBUTION

Early CT literature has attributed gains seen at the receiver to constructive interference (CI) [4]. More recent works [2], [24], in addition to this paper’s analysis in Sect. III, have proposed that, contrary to this assertion, instances of few concurrent transmitters will result in beating (observed as periodic peaks and valleys across a waveform) due to innate CFO inaccuracies between devices defined as RFO in Sect. III. Rather than a CI gain, beating causes periods of *both* constructive *and* destructive interference across the packet, leading to errors during beating valleys, while beating peaks will benefit from a receiver gain. This has recently been demonstrated experimentally by observing the raw in-phase and quadrature (IQ) samples observed when connecting a small number of CT devices to an SDR using coaxial cables [25]. Although these efforts help to better explain some of the processes underpinning CT-based

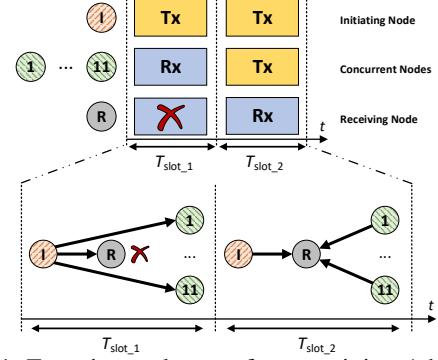


Fig. 4: Experimental setup for examining 1-hop CT.

communication, it is hard to directly witness and evaluate how the occurrence of fundamental physical layer properties affect CT performance. To date, there has been no *over-the-air* testbed experiment able to demonstrate how PHY effects, such as RFO-induced beating and de-synchronization due to clock drift, directly affect the signal observed by a receiving node.

To address this gap, we present experiments that evaluate PHY effects on a 1-hop network of nodes communicating wirelessly by means of CT. Given that PHY properties cause the error probability to flux across the received packet (as shown in Sect. III), we *observe beating by mapping the distribution of bit errors across a packet* when considering a large transmission sample. Specifically, we study the CT performance across the multiple available PHYs supported by the nRF52840 devices in D-Cube. We observe both beating frequencies and de-synchronization effects through analysis of the received error distribution, and demonstrate that their impact on CT performance is highly dependent on the choice of the underlying PHY, on the RFO between transmitting devices, and on the number of concurrent transmitters. We perform all experiments using the *Atomic-SDN* CT stack developed for the EWSN 2019 Dependability Competition [34], [35], [36].

A. Experimental Setup

The D-Cube testbed was configured to provide a single-hop scenario for up to 12 concurrently-transmitting nodes and a fixed receiving node. Fig. 4 demonstrates this setup and shows how the network is able to synchronize all transmitting nodes while limiting packet receptions at the receiver to only those that are a sum of signals from multiple concurrent transmitters. Node R ignores the first transmission from the CT initiator I in T_{slot_1} , while allowing concurrent nodes to receive and synchronize to I . In T_{slot_2} all concurrent nodes synchronously transmit, including the initiator, and are observed at R . Node I was configured to periodically generate and transmit a pseudo-random payload every 250 ms, which was logged before each transmission in T_{slot_1} , alongside an 8-byte CT header. When using IEEE 802.15.4, this was a 119B payload (due to the 127B maximum transmission unit limitation), while a payload of 200B was used for the BLE 5 PHYs⁴. At the receiving node

⁴Although 200B is not the full maximum transmission unit of BLE5, it allows sufficient time to capture wider beating effects while reasonably limiting transmission time when using the BLE 5 125K PHY.




Fig. 5: Errors per packet bit index show a 2kHz beating frequency from a specific CT pair. Low data-rate PHYs experience narrower beating as they span (in time) a greater number of error peaks. The coding used in BLE5 125K handles these errors.




Fig. 6: PRR, PER, and PLR as CT density increases from a single transmitter (no CT) to CT12.

R , the byte arrays of all correctly received and incorrectly received packets were logged in T_{slot_2} . A direct comparison of transmitted and received byte arrays subsequently allows observation of the following four metrics.

- 1) *Bit error distribution* – mapping the distribution of errors across a packet exposes observable periods of gain and interference over time.
- 2) *Packet Reception Ratio (PRR)* – indicating the ability of the physical layer to recover from CT interference.
- 3) *Packet Error Ratio (PER)* – indicating CT interference with a packet after the correct reception of the preamble.
- 4) *Packet Loss Ratio (PLR)* – allowing the observation (by omission) of packets for which there was an energy minimum during a preamble’s reception, resulting in the radio discarding the packet.

This setup was run across the IEEE 802.15.4 and all the BLE 5 PHYs for 2, 3, 4, 6, 8, 10, and 12 concurrently transmitters at -8dBm⁵, where we define CT density as CT2, CT3, ... CT12 respectively. Each experiment was run for ≈ 18 K transmitted packets, representing over 100 total hours of experimental data.

B. Results

Beating effect on different PHYs. Fig. 5 shows the beating frequency for a single CT2 pair. The beating frequency remains consistent across all subfigures but, depending on the underlying PHY rate, the time window spanned by the packet is smaller or larger, resulting in *narrower* or *wider* beating over the same fixed packet length (i.e., high data-rates will suffer from fewer peaks than low data-rates). Taking into account the underlying bit duration and the bit distance between error peaks, it is possible to discern that the beating period for this *specific* CT2 pair is ≈ 0.5 ms (a frequency of 2kHz).

Fig. 7 shows how the PRR and PER are closely linked to the way in which beating manifests across the various PHYs, and supports the findings previously shown through simulation in Fig. 3c (*narrow and strong* beating), where it is likely that this pair experiences *very low noise* (SNR > 25 dB). As discussed in Sect. III, the higher data rate PHYs (BLE 5 2M and 1M)

⁵A transmission power of -8dBm allows to capture sufficient bit errors within a reasonably short time window, minimizing experimentation time.




Fig. 7: PRR and PER for the subplots shown in Fig. 5, where the CT2 pair experiences *significant* beating over a 200B packet.

experience fewer beating valleys. While these uncoded PHYs are unable to recover errors if they fall within a valley, the repetition commonly employed in CT protocols (TX_N = 4 in these experiments) means it is likely that a retransmission will successfully fall between valleys and allow a successful reception of the preamble. This can be observed in the higher error rate for the 1M PHY, which experiences additional beating valleys as opposed to 2M. Furthermore, while it would be natural to assume that the redundancy employed in coded PHYs helps them to better recover from beating errors, our results show that the BLE 5 convolutional coding is unable to cope with significant beating (as seen from the high PER in the BLE 5 500K results shown in Fig. 7). The same applies to the DSSS employed in IEEE 802.15.4, although it helps in recovering errors (despite the higher number of beating valleys caused by the lower data rate). On the other hand, the addition of the Manchester pattern mapper in the BLE 5 125K PHY provides sufficient gain to survive beating.

Increasing CT density. The effect of increasing CT density is explored in Fig. 6. Experiments were run across all PHYs for a single transmitter (*no CT*), as well as increasing CT density from CT2 to CT12. Plots represent an average of multiple experiments run with randomly selected CT forwarders per experiment, while the same pseudo-random forwarding set remains consistent across each PHY. This averaging eliminates bias due to *narrow and strong* beating experienced by CT2 pairs such as Figs. 8e and 8h. Reliability drops at CT3 due to the high data rate of the BLE 5 2M PHY, which requires a large difference in received power between signals to experience the capture effect. This is consistent with recent literature [25] and

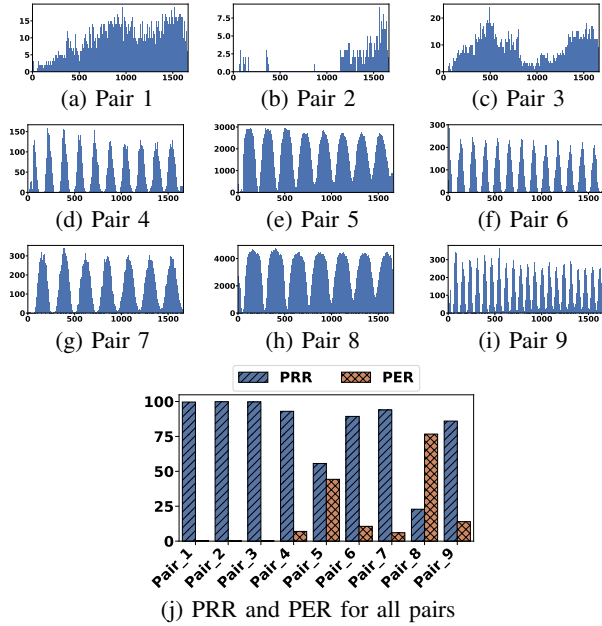


Fig. 8: Error distribution across nine CT2 pairs when using the BLE 5 500K PHY (a–i) and overall CT performance (j).

with our analysis in Sect. III. While the IEEE 802.15.4 PHY still performs well at CT3, its PRR also drops significantly at CT4. Interestingly, the BLE 5 1M PHY shows a gradual drop in performance at mid-level CT densities, while both BLE 5 uncoded PHYs (2M and 1M) experience a PRR ‘rally’ at high CT density. We hypothesize that this is due to the increased diversity (i.e., additional paths and better chance of capturing dominant signals), or to the additional CT converging around an average RFO and spreading the effects of beating.

Beating effect across different CT pairs. Fig. 8 examines the bit error distribution for nine different CT pairs on the BLE 5 500K PHY, chosen due to its beating sensitivity. The absence of the Manchester pattern mapper (as used in the 125K PHY) results in a high degree of corruption across the packet, meaning that beating errors are more prominent. Beating is therefore clearly seen across almost all pairs, with the exception of Figs. 8a and 8b, where the RFO was not significant enough to result in observable beating (i.e., these pairs experience wide beating greater than the packet’s transmission period). Fig. 8j provides additional information about the PRR and PER for each pair, further showing how the RFO between pairs affects the beating width and, consequently, the performance of CT.

C. Key Observations

Recent literature has theorized that the beating effect should have a significant impact on CT performance [25], [19]. The experimental results presented in this section have shown that beating is present in both coded and uncoded BLE 5 PHYs, as well as in the DSSS-based IEEE 802.15.4 PHY. We summarize these results by outlining a number of key observations.

Beating frequencies are device-specific. As shown by Fig. 8, beating frequencies depend on the RFO between device pairs, and one cannot directly extrapolate results from a specific pair.

Preambles are sensitive to beating. While the start of a preamble can randomly coincide with a beating valley or peak, these results are relative to *received* packets (i.e., those for which the preamble was successfully detected) and hence have bias towards a certain initial phase relationship. This bias is further increased by calibration performed by a receiver during the preamble’s reception [30]. As the packet is being received, the beating changes the signal properties and this calibration is no longer optimal; however, the sinusoidal beating frequency will periodically return the signal to that initial calibration. This also explains why beating is visible through an error distribution analysis, and that with no bias (i.e., no preamble) it would present as a flat error distribution.

High data rate PHYs benefit from packet repetition. As shown by Fig. 5, packet transmissions in high data rate PHYs span fewer beating valleys (potentially zero if the packet period is shorter than the beating period). Since the position of peaks and valleys is random, after several repetitions, i.e., with a higher TX_N, it is likely that a packet will not experience a valley during the preamble and will be correctly received. Note that TX_N is a core component of many CT-based protocols.

Low data rate PHYs benefit from the coding gain. Fig. 7 shows a significant difference in reliability between the BLE 5 500K and 125K PHYs. Indeed, BLE 5 500K exhibits the worst performance of any of the PHYs compared in this section. It is likely that the convolutional coding employed by BLE 5 is sensitive to beating errors, while the gain seen using the 125K PHY stems from the additional pattern mapper redundancy over the payload (as mentioned in Sect. III). Similarly, while not achieving the same gains of BLE 5 125K PHY, the DSSS used in IEEE 802.15.4 halves the PER in comparison to the BLE 5 500K PHY. It is worth noting that results in this section do not consider significant external noise or interference, which may particularly penalize the very long packets of the BLE 5 125K PHY, as seen in the following section, and the use of the 500K PHY may again constitute an effective trade-off in harsh environments, particularly to survive intermittent jamming.

V. CT PERFORMANCE OVER DIFFERENT PHYs: EXPERIMENTAL EVALUATION WITH RF INTERFERENCE

This section presents an experimental study on the impact of different PHYs on CT-based protocols in the presence of RF interference. Specifically, we evaluate three CT-based flooding protocols on a large multi-hop network: Glossy [4], Robust Flooding (RoF) [10], and Robust Flooding Single Channel (RoF (SC)), whose operations are depicted in Fig. 9⁶.

The original Glossy [4] approach triggers transmissions after successful receptions, thereby alternating Rx and Tx slots at each hop. However, not only does the Glossy approach operate on a single channel (meaning it is susceptible to RF interference at that frequency) but this reception-triggered Rx-Tx technique means that Rx failures will result in a missed transmission opportunity [10], [37]. Using this technique it is

⁶For a detailed survey on CT-based flooding protocols, please refer to [3].

difficult to resume a CT flood if it is interrupted by interference. An alternative approach was taken by the authors of [10], [38], introducing *back-to-back* transmission slots (i.e. $Rx-Tx-Tx$) alongside robust frequency diversity through per-slot channel hopping. In this Robust Flooding (RoF) approach, the first transmission of a node is still triggered by correct reception, but further transmissions are time-triggered, with nodes synchronously hopping frequency at each slot according to a known offset on a global channel list.

We compare these two approaches (Glossy and RoF) as they are commonly used as primitives to construct more complex CT-based protocols, and are hence representative of wider literature. Furthermore, we introduce a variant of RoF – RoF Single Channel (RoF (SC)) – to observe how this protocol performs w.r.t the single-channel environment used by Glossy.

A. Experimental Setup

We use the D-Cube testbed to evaluate in hardware three key metrics: *end-to-end reliability*, *latency*, and *energy consumption*. For each of the protocols (Glossy, RoF, and RoF (SC)), we consider 3 scenarios characterized by the absence or presence of interference, denoted as *no*, *mild*, and *strong* interference.

D-Cube’s controllable RF interference is generated by observer nodes (Raspberry Pi 3, directly attached to the nRF52840 boards) using JamLab-NG [28]. *Mild interference* (D-Cube level 2) uses a power of 30 mW, generating intermittent (i.e., not continuous) interference for approximately 5 ms every 13 ms period. *Strong interference* (D-Cube level 3) emulates the transmissions of multiple Wi-Fi devices across all the 2.4 GHz band. Each Raspberry Pi 3 node chooses a different channel, generating interference of 200 mW for ≈ 8 ms every 13 ms.

Each protocol is run on D-Cube’s data dissemination scenarios, i.e., sending data from a single source to multiple destinations over a multi-hop network. To emulate event-based scenarios, we configure D-Cube to generate aperiodic messages with short (8B) payload for an alarm scenario and long (64B) payload for a condition monitoring scenario. Other experimental parameters are set as follows. We set the max. number of transmission attempts per node during a flooding period (defined as TX_N and set to three in the example shown in Fig. 9) to 6 for all protocols and fix the flooding periodicity

to 200 ms. In Glossy and RoF (SC) the radio frequency is set to 2.480 GHz (i.e., channel 39 in BLE and channel 26 in IEEE 802.15.4), while RoF hops between 3 different channels (2.4025 GHz, 2.425 GHz, and 2.480 GHz). Finally, we set the transmission power to 0 dBm, which leads to a network diameter between 6 and 10 hops depending on the used layout.

All the results shown in this section utilize publicly available implementations of Glossy⁷ and RoF⁸, which we subsequently ported to the nRF52840-DK platform supported by D-Cube.

B. Results

Reliability. Fig. 10a shows that, under no interference, as data rate increases the evaluated PHYs struggle to maintain reliability. RoF exhibits the highest reliability, while RoF(SC) drops at higher data rates w.r.t. to Glossy. Since the time-triggered transmission approach of RoF and RoF(SC) does not allow nodes to resynchronize at every Rx slot (as in Glossy), nodes can be subject to synchronization errors due to drift. The high data rate PHYs are particularly sensitive to such errors: BLE 5 2M is only able to tolerate CT synchronization errors of up to $0.25 \mu s$ [18]. In general, longer transmission times increase the probability that interference corrupts the packet. This is reflected in the reliability between D-Cube’s long (64B) and short (8B) payloads under all three interference scenarios. Crucially, long transmission times in BLE 5 125K mean it struggles to escape interference, and results in surprisingly poor reliability across all three protocols. We also observe that *back-to-back* repetition of packets in RoF and RoF (SC) improves reliability over Glossy under mild and strong interference. As expected, frequency diversity from RoF’s channel hopping mechanism significantly improves performance of all PHYs – except for the BLE 5 2M. This is likely due to the higher data rate, as interference generated by JamLab-NG is *periodic* across *multiple* channels. High data rate floods are more likely to fall completely within the interference duration, while lower data rates may have transmission slots falling in-between interference periods. The results presented in Fig. 10a also hint that the IEEE 802.15.4 and the BLE 5 500K PHYs exhibit the highest reliability under strong interference when using short and long payloads, respectively.

Latency. The end-to-end latency of CT-based protocols is inherently linked to reliability. Brute-force repetition means that packets *may* successfully be received on poor channels, but much later in the flood. Fig 10b supports this, and we observe significant latency jumps as interference increases. As D-Cube’s latency is only computed based on *received* messages, it is conceivable that low latencies can be achieved even with poor reliability. This is evident under *mild* interference, where the uncoded PHYs exhibit low latency despite poor reliability. In general, coded PHYs (i.e., BLE 5 125K, BLE 5 500K, and IEEE 802.15.4) have better performance w.r.t. latency. Under *no/mild* interference, RoF(SC) BLE 5 500K latency increases in comparison to the other PHYs and protocols, likely due to a

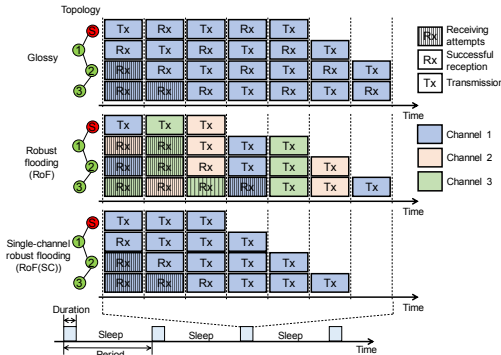


Fig. 9: Operation of Glossy, RoF, and RoF (SC), with the maximum number of transmissions (TX_N) set to three.

⁷<https://sourceforge.net/p/contikiprojects/code/HEAD/tree/ethz.ch/glossy/>

⁸<https://github.com/ETHZ-TEC/robust-flooding - a03f38a>

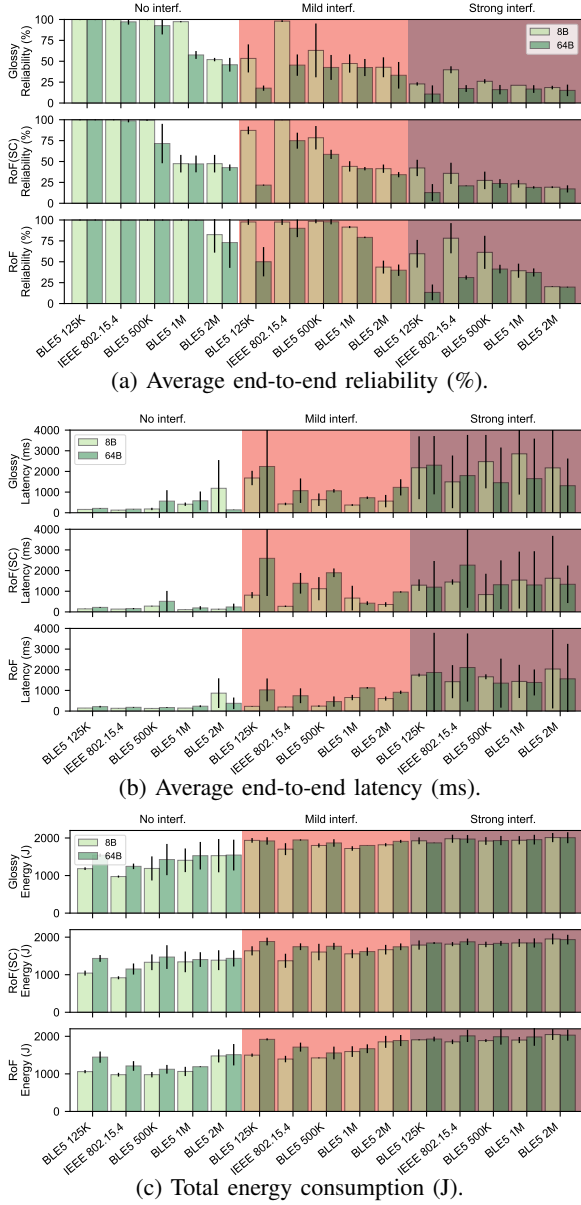


Fig. 10: Average reliability (a), latency (b), and total energy consumption (c) across all D-Cube data dissemination layouts.

combination of the lower data rate, alongside lower reliability due to its beating sensitivity (demonstrated in Sect. III and IV).

Energy. Similarly, node energy is intrinsically linked to reliability, as missed receptions increase radio-on time. Although, in principle, higher data rate PHYs should have a lower energy consumption, this relationship means that for *short* payloads BLE 5 1M and 2M are less energy efficient than the coded PHYs. However, the underlying PHY rate is still a fundamental factor in energy consumption. For *long* payloads the lower data rates of the coded PHYs means that the radio can take a considerable time to transmit, consuming far more energy.

C. Key Observations

Based on these results, we make a number of observations on the network-wide performance of CT-based flooding protocols under RF interference as a function of the employed PHY.

At a network level, high data rate PHYs struggle even in absence of external RF interference. Without the redundancy gains of coded PHYs, high data rate PHYs are sensitive to both de-synchronization and beating at greater CT densities (as per Fig. 6). Even with the added benefit of frequency diversity, RoF’s reliability drops when using the BLE 5 2M PHY.

BLE 5 125K is not necessarily the answer. Although performing well when there is no interference, BLE 5 125K suffers from poor performance as soon as interference kicks in and packet size increases, taking a relatively large hit with respect to latency and reliability in comparison to the other PHYs.

BLE 5 500K and IEEE 802.15.4 perform well under interference. Compared to the other coded PHYs BLE 5 500K achieves higher reliability and similar or lower latency when under interference. However, it performs *worse* in a no interference scenario. This is consistent with the findings in Sect. III, which showed that 500K will outperform other PHYs when there is a significant received power delta, which is likely to occur in high noise conditions. IEEE 802.15.4, on which the majority of CT literature is based, demonstrates similar high reliability under interference, but has a lower data rate. If the level of interference is unknown, CT protocols benefit from transmission on either the BLE 5 500K or IEEE 802.15.4.

RoF’s time-triggered transmission and channel hopping produce significant gains. It is clear from Fig. 10 that the combination of time-triggered transmissions and channel hopping in RoF gives significant gains under all scenarios. However, RoF (SC) shows that without frequency diversity there is a chance that at higher data rates the interference duration may be longer than the flooding period. Increasing TX_N could therefore give greater temporal redundancy and improve protocol reliability under interference.

VI. RELATED WORK

We discuss next related work and highlight how the contributions presented in Sect. III – V advance the state-of-the-art.

CT on different PHYs. After the influential work by Ferrari et al. [4] was published in 2011, a large number of researchers has started to study CT and develop CT-based protocols [3], [5], [6], [7], [8], [9], [39], [40]. While most of the early works targeted exclusively IEEE 802.15.4 devices using the 2.4 GHz band, in the last years, a few studies have shown the feasibility of CT on other physical layers supported by IEEE 802.15.4, such as the UWB PHY [41], [42], [43], as well as sub-GHz short-range [44], [45] and long-range technologies [46], [47].

A number of works have recently focused on studying the feasibility of CT on BLE [18], [19], [25], [48]. Specifically, Al Nahas et al. [18] verified the feasibility of a CT-based flooding protocol, named BlueFlood, on the different BLE 5 PHYs experimentally, and also reported its performance on IEEE 802.15.4 [25]. Schaper [19] studied the conditions to make CT successful in these PHYs in an anechoic chamber.

Different from these studies, our current work does not aim to prove the *feasibility* of CT on different radio technologies or PHYs, but instead to provide an in-depth *characterization* of

Observations on CT Performance	Recommendations
<ul style="list-style-type: none"> • The IEEE 802.15.4 and BLE 5 500K PHYs are effective against external RF interference, but suffer under strong narrow beating, which may cause a significant drop in reliability. • High data rate PHYs help escaping strong narrow beating, but exhibit poor performance in the presence of external RF interference. • The BLE 5 125K PHY is effective against beating, but performs poorly when sending long packets under external RF interference. 	<ul style="list-style-type: none"> • In absence of external RF interference and with a low network density, use BLE 5 2M (or 1M) to ‘widen’ beating and repetitions to exploit temporal redundancy^(*). • In the presence of external RF interference, use BLE 5 125K only for short packets. Consider this PHY also to escape beating^(*). • In the presence of strong external RF interference, use IEEE 802.15.4 for shorter packets and BLE 5 500K for longer packets.

PHY	Beating Errors		External RF Interference	
	short packet	long packet	short packet	long packet
BLE 5 125K	↑	↗	↗	↓
BLE 5 500K	↘	↓	↗	↑
BLE 5 1M	↗	↘	↘	↘
BLE 5 2M	↑	↗	↓	↓
IEEE 802.15.4	↘	↓	↑	↗

^(*)The choice of PHY to cope with beating should also be made based on the application’s latency, energy, and RF range requirements.

TABLE I: Left table: summary of observations on CT performance over different PHYs and corresponding recommendations. Right table: suitability of the different PHYs to mitigate the presence of beating errors and external RF interference. Up arrows indicate a higher suitability and better performance; down arrows suggest a lower suitability and poor performance.

the role of the physical layer on the reliability and efficiency of CT-based solutions employing IEEE 802.15.4 and BLE 5. To the best of our knowledge, we are the first to do this in a systematic manner by demonstrating experimentally the impact of errors induced by de-synchronization and beating distortion in CT-based protocols as a function of the employed PHY.

CT performance under interference. Several works have shown that CT-based data collection and dissemination protocols can outperform conventional routing-based solutions in terms of reliability, end-to-end, and energy consumption even in the presence of harsh radio interference [10], [11], [12], [13], as also highlighted in the context of the EWSN dependability competition series [14], [26], [27]. To sustain a dependable performance under interference, CT-based solutions have been enriched, among others, with mechanisms such as local opportunistic retransmissions [4], [49], [50], channel-hopping [10], [11], [51], [52], [53], network coding [8], [9], [54], [55], [56], noise detection [11], [12], stretched preambles [13], data freezing [12], as well as an improved understanding of the network state [57], [58], [59].

However, most of these protocols have been implemented for and evaluated with IEEE 802.15.4 technology only. In this paper, we are the first to study the performance of CT-based data collection protocols on a large scale under interference *as a function of the employed PHY*. We did this by evaluating the dependability of CT-based protocols on a large scale using modern multi-radio platforms supporting several PHYs, and by analyzing the impact of other factors such as the length of the transmitted messages and the harshness of the interference.

Impact of beating effect on CT. A few studies have tried to underpin the foundations of concurrent transmissions on a signal level. Ferrari et al. [4] have simulated CT signals with Matlab and explained how accurately packets should be aligned in order to design reliable protocols. Other works [60], [61], [62] have analyzed CT signals theoretically and argued that it is difficult to generate ideal constructive interference, due to the timing errors caused by radio propagation and clock drift.

Liao et al. have been the first to argue that there exists a beating effect caused by innate CFO between device oscillators [2]. Specifically, in [2], the resultant signals are generated by Matlab and a TelosB node was used to observe how the DSSS modulation in IEEE 802.15.4 saves CT signals from the beating effect. More recent studies have demonstrated

these beating effects generate periods of both constructive and destructive interference by observing the raw IQ samples of devices connected to an SDR using coaxial cables [25].

In this paper we are the first to demonstrate how PHY effects such as CFO-induced beating, and de-synchronization from hardware clock drift, directly affect the signal observed by a receiving node *using over-the-air testbed experiments*.

VII. DISCUSSION AND FUTURE WORK

This paper provides the first systematic experimental evaluation about the impact of the physical layer on CT performance. While these results are specific to the IEEE 802.15.4 and BLE 5 PHYs, they provide important insights on how the choice of the PHY can exacerbate or reduce errors due to beating, de-synchronization, and external RF interference.

Our main observations and recommendations are summarized in Table I. With modern chipsets now supporting real-time PHY switching [22], [23], these insights can guide further research into CT protocols that take advantage of multi-PHY capabilities. For example, future protocols may take advantage of repetitions in uncoded PHYs to combat beating, while mitigating RF interference through coded PHYs.

Yet, while these findings are important for the design of CT protocols, there are a number of key areas that require additional research. Foremost, analytical and mathematical modeling would help to substantiate the observations presented in Sect. III and IV. Crucially, a greater understanding is needed around how beating errors affect a protocol’s scalability on a network level, as real-world RF conditions make it difficult to determine how beating will manifest at each node. Furthermore, as CFO is sensitive to temperature, the RFO may change over time resulting in changes to the beating frequency. Using the technique outlined in this paper would allow continuous monitoring of the beating. Finally, techniques such as *interleaving*, i.e., bit shuffling, (which improves the robustness of FEC with respect to burst errors) and *bit voting* are missing in the analyzed physical layers, and would be a effective addition to increase the reliability of CT.

ACKNOWLEDGEMENTS

This work was partially supported by the National Natural Science Foundation of China under Grant 61902188, and in part by the China Postdoctoral Science Foundation under Grant 2020T130304. This work has also been partly performed within the LEAD project “Dependable Internet of Things in Adverse Environments” funded by Graz University of Technology.

REFERENCES

- [1] Krijn Leentvaar and Jan H. Flint. The Capture Effect in FM Receivers. *IEEE Transactions on Communications*, 24(5):531–539, May 1976.
- [2] Chun-Hao Liao, Yuki Katsumata, Makoto Suzuki, and Hiroyuki Morikawa. Revisiting the So-Called Constructive Interference in Concurrent Transmission. In *Proceedings of the 41st International Conference on Local Computer Networks (LCN)*, pages 280–288. IEEE, November 2016.
- [3] Marco Zimmerling, Luca Mottola, and Silvia Santini. Synchronous Transmissions in Low-Power Wireless: A Survey of Communication Protocols and Network Services. *CORR – arXiv preprint 2001.08557*, January 2020.
- [4] Federico Ferrari, Marco Zimmerling, Lothar Thiele, and Olga Saukh. Efficient Network Flooding and Time Synchronization with Glossy. In *Proceedings of the 10th International Conference on Information Processing in Sensor Networks (IPSN)*, pages 73–84. IEEE, April 2011.
- [5] Federico Ferrari, Marco Zimmerling, Luca Mottola, and Lothar Thiele. Low-Power Wireless Bus. In *Proceedings of the 10th International Conference on Embedded Network Sensor Systems (SenSys)*, pages 1–14. ACM, November 2012.
- [6] Makoto Suzuki, Yasuta Yamashita, and Hiroyuki Morikawa. Low-Power, End-to-End Reliable Collection Using Glossy for Wireless Sensor Networks. In *Proceedings of the 77th International Vehicular Technology Conference (VTC)*, pages 1–5. IEEE, June 2013.
- [7] Timofei Istomin, Amy Lynn Murphy, Gian Pietro Picco, and Usman Raza. Data Prediction + Synchronous Transmissions = Ultra-Low Power Wireless Sensor Networks. In *Proceedings of the 14th International Conference on Embedded Network Sensor Systems (SenSys)*, pages 83–95. ACM, November 2016.
- [8] Manjunath Doddavenkatappa, Mun Choon Chan, and Ben Leong. Splash: Fast Data Dissemination with Constructive Interference in Wireless Sensor Networks. In *Proceedings of the 10th International Symposium on Networked Systems Design and Implementation (NSDI)*, pages 269–282. USENIX, April 2013.
- [9] Wan Du, Jansen Christian Liando, Huanle Zhang, and Mo Li. Pando: Fountain-Enabled Fast Data Dissemination With Constructive Interference. *IEEE/ACM Transactions on Networking*, 25(2):820–833, April 2017.
- [10] Roman Lim, Reto Da Forno, Felix Sutton, and Lothar Thiele. Competition: Robust Flooding using Back-to-Back Synchronous Transmissions with Channel-Hopping. In *Proceedings of the 14th International Conference on Embedded Wireless Systems and Networks (EWSN), competition session*, pages 270–271. Junction Publishing, February 2017.
- [11] Timofei Istomin, Matteo Trobinger, Amy Lynn Murphy, and Gian Pietro Picco. Interference-Resilient Ultra-low Power Aperiodic Data Collection. In *Proceedings of the 17th International Conference on Information Processing in Sensor Networks (IPSN)*, pages 84–95. ACM, April 2018.
- [12] Xiaoyuan Ma, Peilin Zhang, Ye Liu, Carlo Alberto Boano, Hyung-Sin Kim, Jianming Wei, and Jun Huang. Harmony: Saving Concurrent Transmissions from Harsh RF Interference. In *Proceedings of the 39th International Conference on Computer Communication (INFOCOM)*. IEEE, July 2020.
- [13] Antonio Escobar-Molero, Javier Garcia-Jimenez, Jirka Klaue, Fernando Moreno-Cruz, Borja Saez, Francisco J Cruz, Unai Ruiz, and Angel Corona. Competition: RedNodeBus, Stretching out the Preamble. In *Proceedings of the 16th International Conference on Embedded Wireless Systems and Networks (EWSN), competition session*, pages 304–305. Junction Publishing, February 2019.
- [14] Carlo Alberto Boano, Markus Schuß, and Kay Römer. EWSN Dependability Competition: Experiences and Lessons Learned. *IEEE Internet of Things Newsletter*, March 2017.
- [15] Joseph Polastre, Robert Szewczyk, and David E. Culler. Telos: Enabling Ultra-Low Power Wireless Research. In *Proceedings of the 4th International Symposium on Information Processing in Sensor Networks (IPSN)*, pages 364–369. IEEE, April 2005.
- [16] IEEE 802.15.4 Working Group. *IEEE Standard for Low-Rate Wireless Networks*, IEEE Std 802.15.4-2015 (Revision of IEEE Std 802.15.4-2011, IEEE Std 802.15.4-2006, and IEEE Std 802.15.4-2003) edition, April 2016.
- [17] Matthias Wilhelm, Vincent Lenders, and Jens B. Schmitt. On the Reception of Concurrent Transmissions in Wireless Sensor Networks. *IEEE Transactions on Wireless Communications*, 13(12):6756–6767, August 2014.
- [18] Beshr Al Nahas, Simon Duquennoy, and Olaf Landsiedel. Concurrent Transmissions for Multi-Hop Bluetooth 5. In *Proceedings of the 16th International Conference on Embedded Wireless Systems and Networks (EWSN)*, pages 130–141. Junction Publishing, February 2019.
- [19] Anna-Brit Schaper. Truth be Told: Benchmarking BLE and IEEE 802.15.4. Master’s thesis, ETH Zurich, Zurich, Switzerland, November 2019.
- [20] Romain Jacob, Anna-Brit Schaper, Andreas Biri, Reto da Forno, and Lothar Thiele. Synchronous Transmissions on Bluetooth 5 and IEEE 802.15.4: A Replication Study. In *Proceedings of the 3rd International Workshop on Benchmarking Cyber-Physical Systems and Internet of Things (CPS-IoTBench)*, September 2020.
- [21] Michael Spörk, Carlo Alberto Boano, and Kay Römer. Performance and Trade-offs of the new PHY Modes of BLE 5. In *Proceedings of the International Workshop on Pervasive Systems in the IoT Era (PERSIST-IoT)*, pages 7–12. ACM, July 2019.
- [22] Texas Instruments. CC2652R SimpleLink Multiprotocol 2.4 GHz Wireless MCU datasheet, Rev. G. [Online] <https://www.ti.com/product/CC2652R> – Last accessed: 2020-05-26.
- [23] Nordic Semiconductors. nRF52840 Product Specification, v1.1. [Online] https://infocenter.nordicsemi.com/pdf/nRF52840_PS_v1.1.pdf – Last accessed: 2020-05-26.
- [24] Antonio Escobar-Molero. Improving Reliability and Latency of Wireless Sensor Networks Using Concurrent Transmissions. *at-Automatisierungstechnik*, 67(1):42–50, 2019.
- [25] Beshr Al Nahas, Antonio Escobar-Molero, Jirka Klaue, Simon Duquennoy, and Olaf Landsiedel. BlueFlood: Concurrent Transmissions for Multi-Hop Bluetooth 5—Modeling and Evaluation. *CORR – arXiv preprint 2002.12906*, February 2020.
- [26] Markus Schuß, Carlo Alberto Boano, Manuel Weber, and Kay Römer. A Competition to Push the Dependability of Low-Power Wireless Protocols to the Edge. In *Proceedings of the 14th International Conference on Embedded Wireless Systems and Networks (EWSN)*, pages 54–65. Junction Publishing, February 2017.
- [27] Markus Schuß, Carlo Alberto Boano, and Kay Römer. Moving Beyond Competitions: Extending D-Cube to Seamlessly Benchmark Low-Power Wireless Systems. In *Proceedings of the 1st International Workshop on Benchmarking Cyber-Physical Networks and Systems (CPSBench)*, pages 30–35. IEEE, April 2018.
- [28] Markus Schuß, Carlo Alberto Boano, Manuel Weber, Matthias Schulz, Matthias Hollick, and Kay Römer. JamLab-NG: Benchmarking Low-Power Wireless Protocols under Controllable and Repeatable Wi-Fi Interference. In *Proceedings of the 16th International Conference on Embedded Wireless Systems and Networks (EWSN)*, pages 83–94. Junction Publishing, February 2019.
- [29] Subbarayan Pasupathy. Minimum Shift Keying: A Spectrally Efficient Modulation. *IEEE Communications Magazine*, 17(4):14–22, July 1979.
- [30] Bluetooth Working Group. Bluetooth Core Specification, Revision 5.2, December 2019.
- [31] Antonio Escobar-Molero. *Using Concurrent Transmissions to Improve the Reliability and Latency of Low-Power Wireless Mesh Networks*. Phd thesis, RWTH Aachen University, Aachen, Germany, June 2020.
- [32] Andrew J. Viterbi. Error Bounds for Convolutional Codes and an Asymptotically Optimum Decoding Algorithm. *IEEE Transactions on Information Theory*, 13(2):260–269, April 1967.
- [33] Bluetooth Blog. Exploring Bluetooth 5 – Going the Distance. [Online] <https://www.bluetooth.com/blog/exploring-bluetooth-5-going-the-distance/> – Last accessed: 2020-05-26.
- [34] Michael Baddeley, Usman Raza, Mahesh Sooriyabandara, George Oikonomou, Reza Nejabati, and Dimitra Simeonidou. Atomic-SDN: Is Synchronous Flooding the Solution to Software-Defined Networking in IoT? *IEEE Access*, 7:96019–96034, May 2019.
- [35] Michael Baddeley. *Software Defined Networking for the Industrial Internet of Things*. PhD thesis, Dept. of Electrical and Electronic Engineering, Univ. Bristol, UK, 2020.
- [36] Michael Baddeley, Aleksandar Stanoev, Usman Raza, Yichao Jin, and Mahesh Sooriyabandara. Competition: Adaptive Software Defined Scheduling of Low Power Wireless Networks. In *Proceedings of the 16th International Conference on Embedded Wireless Systems and Networks (EWSN), competition session*, pages 298–299. Junction Publishing, February 2019.
- [37] Xiaoyuan Ma, Peilin Zhang, Weisheng Tang, Xin Li, Wangji He, Fuping Zhang, Jianming Wei, and Oliver Theel. Using Enhanced OFFCOIN to Monitor Multiple Concurrent Events under Adverse Conditions. In *Proceedings of the 15th International Conference on Embedded Wireless*

- Systems and Networks (EWSN), competition session*, pages 211–212. Junction Publishing, February 2018.
- [38] Usman Raza, Yichao Jin, and Mahesh Sooriyabandara. Competition: Synchronous Transmissions based Flooding for Dependable Internet of Things. In *Proceedings of the 14th International Conference on Embedded Wireless Systems and Networks (EWSN), competition session*, pages 278–279. Junction Publishing, February 2017.
 - [39] Olaf Landsiedel, Federico Ferrari, and Marco Zimmerling. Chaos: Versatile and efficient all-to-all data sharing and in-network processing at scale. In *Proceedings of the 11th International Conference on Embedded Networked Sensor Systems (SenSys)*, pages 1–14. ACM, November 2013.
 - [40] Tengfei Chang, Thomas Watteyne, and Xavier Vilajosana Pedro Henrique Gomes. Constructive Interference in 802.15.4: A Tutorial. *IEEE Communications Surveys and Tutorials*, 21(1):217–237, September 2018.
 - [41] Benjamin Kempke, Pat Pannuto, Bradford Campbell, and Prabal Dutta. SurePoint: Exploiting Ultra Wideband Flooding and Diversity to Provide Robust, Scalable, High-Fidelity Indoor Localization. In *Proceedings of the 14th ACM International Conference on Embedded Network Sensor Systems (SenSys)*, pages 137–149. ACM, November 2016.
 - [42] Davide Vecchia, Pablo Corbalán, Timofei Istomin, and Gian Pietro Picco. Playing with Fire: Exploring Concurrent Transmissions in Ultra-wideband Radios. In *Proceedings of the 18th International Conference on Sensing, Communication and Networking (SECON)*, pages 1–9. IEEE, June 2019.
 - [43] Diego Lobba, Matteo Trobinger, Davide Vecchia, Timofei Istomin, and Gian Pietro Picco. Concurrent Transmissions for Multi-hop Communication on Ultra-wideband Radios. In *Proceedings of the 17th International Conference on Embedded Wireless Systems and Networks (EWSN)*, pages 132–143. Junction Publishing, February 2020.
 - [44] Chun-Hao Liao, Makoto Suzuki, and Hiroyuki Morikawa. Toward Robust Concurrent Transmission for Sub-GHz Non-DSSS Communication. In *Proceedings of the 14th International Conference on Embedded Network Sensor Systems (SenSys), poster session*, pages 354–355. ACM, November 2016.
 - [45] Jan Beutel, Roman Trüb, Reto Da Forno, Markus Wegmann, Tonia Gsell, Romain Jacob, Michael Keller, Felix Sutton, and Lothar Thiele. The Dual Processor Platform Architecture. In *Proceedings of the 18th International Conference on Information Processing in Sensor Networks (IPSN), demo session*, pages 335–336. IEEE, April 2019.
 - [46] Chun-Hao Liao, Guibing Zhu, Daiki Kuwabara, Makoto Suzuki, and Hiroyuki Morikawa. Multi-Hop LoRa Networks Enabled by Concurrent Transmission. *IEEE Access*, 5:21430–21446, September 2017.
 - [47] Xiaoyuan Ma, Dan Li, Fengxu Yang, Carlo Alberto Boano, Pei Tian, and Jianming Wei. Chirpbox – A Low-Cost LoRa Testbed Solution. In *Proceedings of the 17th International Conference on Embedded Wireless Systems and Networks (EWSN), poster session*. Junction Publishing, February 2020.
 - [48] Coen Roest. Enabling the Chaos Networking Primitive on Bluetooth LE. Master’s thesis, TU Delft, Delft, The Netherlands, October 2015.
 - [49] Felix Sutton, Bernhard Buchli, Jan Beutel, and Lothar Thiele. Zippy: On-Demand Network Flooding. In *Proceedings of the 13th International Conference on Embedded Networked Sensor Systems (SenSys)*, pages 45–58. ACM, November 2015.
 - [50] Makoto Suzuki, Chun-Hao Liao, Sotaro Ohara, Kyoichi Jinno, and Hiroyuki Morikawa. Wireless-Transparent Sensing. In *Proceedings of the 14th International Conference on Embedded Wireless Systems and Networks (EWSN)*, pages 66–67. Junction Publishing, February 2017.
 - [51] Philipp Sommer and Yvonne-Anne Pignolet. Dependable Network Flooding Using Glossy with Channel-Hopping. In *Proceedings of the 17th International Conference on Embedded Wireless Systems and Networks (EWSN), competition session*, page 303. Junction Publishing, February 2016.
 - [52] Antonio Escobar-Molero, Francisco J Cruz, Javier Garcia-Jimenez, Jirka Klauke, and Angel Corona. RedFixHop with Channel Hopping: Reliable Ultra-low-latency Network Flooding. In *Proceedings of the 16th International Conference on Design of Circuits and Integrated Systems (DCIS)*, pages 1–4. IEEE, November 2016.
 - [53] Manjunath Doddavenkatappa and Mun Choon Chan. P3: A Practical Packet Pipeline using Synchronous Transmissions for Wireless Sensor Networks. In *Proceedings of the 13th International Conference on Information Processing in Sensor Networks (IPSN)*, pages 203–214. IEEE, April 2014.
 - [54] Dingwen Yuan and Matthias Hollick. Ripple: High-throughput, Reliable and Energy-efficient Network Flooding in Wireless Sensor Networks. In *Proceedings of the 16th International Symposium on A World of Wireless, Mobile and Multimedia Networks (WoWMoM)*, pages 1–9. IEEE, June 2015.
 - [55] Mobashir Mohammad and Mun Choon Chan. Codecast: Supporting Data Driven In-Network Processing for Low-Power Wireless Sensor Networks. In *Proceedings of the 17th International Conference on Information Processing in Sensor Networks (IPSN)*, pages 72–83. IEEE, April 2018.
 - [56] Carsten Herrmann, Fabian Mager, and Marco Zimmerling. Mixer: Efficient Many-to-All Broadcast in Dynamic Wireless Mesh Networks. In *Proceedings of the 16th International Conference on Embedded Network Sensor Systems (SenSys)*, pages 145–158. ACM, November 2018.
 - [57] Doug Carlson, Marcus Chang, Andreas Terzis, Yin Chen, and Omprakash Gnawali. Forwarder Selection in Multi-transmitter Networks. In *Proceedings of the 18th International Conference on Distributed Computing in Sensor Systems (DCOSS)*, pages 1–10. IEEE, May 2013.
 - [58] Martina Brachmann, Olaf Landsiedel, and Silvia Santini. Concurrent Transmissions for Communication Protocols in the Internet of Things. In *Proceedings of the 41st International Conference on Local Computer Networks (LCN)*, pages 406–414. IEEE, November 2016.
 - [59] Chayan Sarkar, R. Venkatesha Prasad, Raj Thilak Rajan, and Koen Langendoen. Sleeping Beauty: Efficient Communication for Node Scheduling. In *Proceedings of the 13th International Conference on Mobile Ad Hoc and Sensor Systems (MASS)*, pages 56–64. IEEE, October 2016.
 - [60] Vijay S. Rao, Madhusudan Koppal, RangaRao Venkatesha Prasad, T.V. Prabhakar, Chayan Sarkar, and Ignas Niemegeers. Murphy Loves CI: Unfolding and Improving Constructive Interference in WSNs. In *Proceedings of the 35th International Conference on Computer Communications (INFOCOM)*, pages 1–9. IEEE, April 2016.
 - [61] Yin Wang, Yuan He, Dapeng Cheng, Yunhao Liu, and Xiang-yang Li. TriggerCast: Enabling Wireless Constructive Collisions. In *Proceedings of the 32th International Conference on Computer Communications (INFOCOM)*, pages 480–484. IEEE, April 2013.
 - [62] Yin Wang, Yunhao Liu, Yuan He, Xiang-Yang Li, and Dapeng Cheng. Disco: Improving Packet Delivery via Deliberate Synchronized Constructive Interference. *IEEE Transactions on Parallel and Distributed Systems*, 26(3):713–723, March 2015.