

Understanding Concurrent Transmissions

The Impact of Carrier Frequency Offset and RF Interference on Physical Layer Performance

MICHAEL BADDELEY, Secure Systems Research Center, Technology Innovation Institute, UAE
 CARLO ALBERTO BOANO, Institute of Technical Informatics, Graz University of Technology, Austria
 ANTONIO ESCOBAR-MOLERO, RedNodeLabs UG, Germany
 YE LIU, College of Artificial Intelligence, Nanjing Agricultural University, China
 XIAOYUAN MA, SKF Group, China
 VICTOR MAROT, University of Bristol, United Kingdom
 USMAN RAZA, Waymap Ltd., United Kingdom
 KAY RÖMER, Institute of Technical Informatics, Graz University of Technology, Austria
 MARKUS SCHUSS, Institute of Technical Informatics, Graz University of Technology, Austria
 ALEKSANDAR STANOEVI, Bristol Research and Innovation Lab, Toshiba Europe Ltd., United Kingdom

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 an 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 relative carrier frequency offsets on the performance of CT highly depends on the choice of the underlying physical layer. We then confirm these observations experimentally on real hardware and with varying environmental conditions 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 data collection and dissemination protocols in the presence of RF interference on a large-scale testbed, deriving insights on how the employed physical layer affects their dependability.

CCS Concepts: • **Networks** → *Network performance evaluation; Network protocols*; • **Computer systems organization** → *Embedded and cyber-physical systems; Dependable and fault-tolerant systems and networks; Reliability; Sensor networks*.

Additional Key Words and Phrases: Beating effect, Concurrent transmissions, Data collection, Data dissemination, D-Cube, Dependability, IoT, Low-power wireless, nRF52840, Physical layer, Protocols, Synchronous flooding, Temperature, Testbeds.

Authors' addresses: Michael Baddeley, Secure Systems Research Center, Technology Innovation Institute, Mazdar City, Abu Dhabi, UAE, P.O. Box 9639, michael.baddeley@tii.ae; Carlo Alberto Boano, Institute of Technical Informatics, Graz University of Technology, Inffeldgasse 16, Graz, Steiermark, Austria, 8010, cboano@tugraz.at; Antonio Escobar-Molero, RedNodeLabs UG, Ickstattstr. 18, Munich, Germany, 80469, antonio@rednodelabs.com; Ye Liu, College of Artificial Intelligence, Nanjing Agricultural University, No. 40, Dian Jiang Tai Road, Nanjing, China, 210031, yeliu@njau.edu.cn; Xiaoyuan Ma, SKF Group, No. 1189, Yuan Qi Road, Shanghai, China, 201814, ma.xiaoyuan.mail@gmail.com; Victor Marot, University of Bristol, Queens Road, Bristol, United Kingdom, BS81QU, hk19236@bristol.ac.uk; Usman Raza, Waymap Ltd., Dowgate Hill House, London, United Kingdom, EC4R 2SU, usman.raza@waymapnav.org; Kay Römer, Institute of Technical Informatics, Graz University of Technology, Inffeldgasse 16, Graz, Steiermark, Austria, 8010, roemer@tugraz.at; Markus Schuss, Institute of Technical Informatics, Graz University of Technology, Inffeldgasse 16, Graz, Steiermark, Austria, 8010, markus.schuss@tugraz.at; Aleksandar Stanoev, Bristol Research and Innovation Lab, Toshiba Europe Ltd., 32 Queen Square, Bristol, United Kingdom, BS1 4ND, Aleksandar.Stanoev@toshiba-bril.com.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, or post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

© 2023 Copyright held by the owner/author(s). Publication rights licensed to ACM.

1550-4859/2023/6-ART1 \$15.00

<https://doi.org/10.1145/3604430>

ACM Reference Format:

Michael Baddeley, Carlo Alberto Boano, Antonio Escobar-Molero, Ye Liu, Xiaoyuan Ma, Victor Marot, Usman Raza, Kay Römer, Markus Schuß, and Aleksandar Stanoev. 2023. **Understanding Concurrent Transmissions: The Impact of Carrier Frequency Offset and RF Interference on Physical Layer Performance**. *ACM Trans. Sensor Netw.* 1, 1, Article 1 (June 2023), 38 pages. <https://doi.org/10.1145/3604430>

1 INTRODUCTION

Throughout the last decade, the development of communication protocols based on Concurrent Transmissions (CT) has driven a breakthrough in highly-performant and dependable low-power wireless networking protocols¹. CT-based protocols intentionally let multiple relaying nodes forward packets by simultaneously broadcasting them on the same carrier frequency. Thanks to the capture effect [44] and to non-destructive interference [45], nodes overhearing these CT 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 [94].

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 [29]. 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 hence avoid the overhead of routing and link-based communication [94].

A large body of work has proposed CT-based data collection [28, 37, 80] and dissemination [19, 21, 22] 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 [27, 38, 48, 50], as shown by four editions of the EWSN dependability competition [11].

However, the vast majority of these protocols have only been implemented and verified experimentally using off-the-shelf platforms based on IEEE 802.15.4 radios operating in the 2.4 GHz ISM band (e.g., the popular, but rather outdated TelosB mote [60]). 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 [36], where DSSS provides the coding robustness needed by CT to be sufficiently reliable [89].

The rise of multi-PHY IoT platforms and the role of the physical layer. An experimental study by Al Nahas et al. [2] 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 [39, 66]. Indeed, the most recent version of Bluetooth Low Energy (BLE 5) supports four PHYs that largely differ in terms of data rate and robustness [73]: 2M (2 Mbps), which doubles the nominal throughput of the original 1M PHY (1 Mbps), 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 Texas Instruments' CC2652R [81], the Silicon Labs' EFR32MG [71], the STMicroelectronics' STM32WB [77], the NXP Semiconductors' K32W061 [57], as well as the Nordic Semiconductors' nRF52840 [55] and nRF5340 [56], which all support 2.4 GHz IEEE 802.15.4 OQPSK-DSSS alongside the four BLE 5 PHYs².

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

¹Different terminologies exist in the literature to refer to this concept, including *synchronous* transmissions [94], *simultaneous* transmissions [18], and *concurrent* transmissions [89]. We will stick to the latter in the remainder of this paper.

²Especially the nRF52840 has become popular in the research community in recent years, as a number of researchers started to use it to experimentally validate their low-power wireless solutions [2, 4, 8, 30, 32, 39, 43, 74–76].

slightly different frequencies, introduced in Sect. 2) induced by relative carrier frequency offset [45], and de-synchronization due to clock drift [24] affect the reliability of the received signal in the presence of multiple concurrent transmitters. Importantly, as a function of underlying carrier frequency offsets, the use of non temperature controlled oscillators in low-power radios ensures that beating and clock drift will respond in-line with environmental (both indoor and outdoor) temperature fluctuations - and therefore the impact of these effects may change over time. Furthermore, there is no experimental work studying whether the performance of CT in harsh RF environments (i.e., in the presence of RF interference) differs depending on the underlying PHY.

Shedding light on these aspects is important to (i) provide a better understanding about 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) and low level access to the received baseband signal.

Secondly, while recent selected works have demonstrated the feasibility of CT over BLE in single-hop experiments [2, 39, 53, 66], 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* [55, 81], 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 IEEE 802.15.4 PHY and 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 [67, 69], recently enhanced with 50 nRF52840-DK devices and with the ability of controlling the temperature of the nodes [68], to observe both beating frequencies and de-synchronization effects on real hardware through an 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, on environmental temperature fluctuations, 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 retransmissions; (ii) if the power delta between signals is insufficient, then BLE 5 convolutional coding does not sustain reliable CT; (iii) the pattern mapper used in BLE 5 125K allows it to effectively handle narrow beating. Among others, our experimental evaluation shows that: (iv) observed beating frequencies at the receiver can be quantified by applying a Fast Fourier Transform (FFT) over a histogram of received bit errors; (v) beating frequencies are highly sensitive to the temperature at transmitting nodes, in line with the temperature response of the RF oscillator; (vi) with steady temperature at the transmitters, the beating frequency will remain consistent over time.

We further use D-Cube to perform the first experimental study on the performance of different CT-based data collection and dissemination protocols as a function of the underlying PHY in the presence of RF interference on a large scale. To this end, we use D-Cube's JamLab-NG functionality [70] to generate artificial Wi-Fi interference and stress-test the performance of CT-based data dissemination protocols such as Glossy [29] and robust flooding (RoF) [48] under *no*, *mild*, and *strong* interference. We then repeat the same study for data collection protocols




Fig. 1. **Illustration of the beating effect.** Concurrent transmitters (S1 and S2) experience a mismatch between the frequencies of their carrier signals due to imperfect crystal oscillators. As a consequence, the concurrent signals combine to produce periods of *constructive* and *destructive* interference (denoted as CI and DI, respectively) at the receiver (R).

such as Crystal [37] and its multi-channel and interference-resilient extension [38]. Our results allow us to derive important insights on which PHYs are effective to help CT-based protocols in mitigating the impact of RF interference. Such insights apply to both data collection and dissemination protocols, and 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 and on the beating effect in Sect. 2, this paper makes the following specific contributions:

- We simulate the performance of CT for the IEEE 802.15.4 PHY and all BLE 5 PHYs, highlighting the strong impact of beating on the resulting packet error rate (Sect. 3).
- We are the first to experimentally observe beating frequencies and de-synchronization effects on real hardware and wireless channel for different PHYs through a novel technique whereby analysis of the bit error distribution across received packets exposes the underlying beating frequency (Sect. 4).
- We demonstrate that the beating frequency can be quantified by applying a Fast Fourier Transform (FFT) across the received bit error distribution, potentially allowing continuous and real-time monitoring of beating in future CT-based systems (Sect. 4).
- We show that the sensitivity of radio oscillators to temperature variations impacts the relative carrier frequency offset between devices and consequently the CT-induced beating frequency at the receiver (Sect. 4).
- We experimentally evaluate the performance of CT-based data collection and dissemination protocols in the presence of RF interference, and provide insights on how the employed PHYs affect their reliability, end-to-end latency, as well as energy efficiency (Sect. 5).

We then describe related work in Sect. 6 to highlight how our insights align with existing literature, and conclude our paper in Sect. 7, along with a discussion on future work.

2 DEMYSTIFYING CONCURRENT TRANSMISSIONS IN LOW-POWER WIRELESS NETWORKS

Humans, as well as wireless devices, typically communicate with each other *politely*, meaning that basic principles of conversational manners and wireless design favored strategies that prevent multiple speakers/transmitters to talk over each other. People's tendency to listen before talking and not interrupting others did inspire wireless MAC designs such as carrier sense multiple access with collision avoidance (CSMA/CA) when more intrusive approaches like ALOHA had proven to achieve less effective communication throughput. Time Division Multiple Access (TDMA), the other dominant communication paradigm, emulates the turn-taking in human conversation.

CT-based communication is, however, a more disruptive approach and is beyond the polite conversational conventions mentioned above. CT encourages multiple wireless devices to transmit data *exactly at the same time*. Finding analogies of CT in a normal human situation is not difficult though. Think about listening, understanding, and potentially remembering an unheard song sung by multiple people at the same time in a pub. The ability to successfully decode data transmitted by multiple concurrent transmitters was pioneered by Glossy [29] a decade ago for low-power wireless networks. In Glossy, all the transmitters send the same data just like the same song is sung by all the pub attendees. Now, imagine listening to an argument between a group of people who are talking over each other, trying to get their *different* point of views across. Yet, the loudest voice and often its contents can be understood even in this seemingly chaotic situation. Low power wireless radios, not very different than human listeners, can understand the loudest of many signals due to the *capture effect* of radios [44, 88]³. This phenomenon enabled new flexible designs that permit multiple concurrent transmitters to send different data packets simultaneously. Chaos [42] and Crystal [37] are two well-known examples of CT-based protocols building on top of Glossy, but the literature on CT-based protocols is constantly growing [94].

Low power CT-based protocols are tested with several wireless technologies, including different variations of IEEE 802.15.4 (narrowband/ultra-wideband, 2.4 GHz/sub-GHz), BLE, LoRa, etc. In fact, diverse physical models which are based on SINR have been proposed and taken as an important factor when designing CT protocols [17]. Most experimental studies, however, are based on 2.4 GHz-based short range technologies, which include the narrowband IEEE 802.15.4 and BLE physical layers. The narrowband IEEE 802.15.4 specifications use OQPSK modulation with half-sine pulse shaping in the 2.4 GHz band. Many commercial off-the-shelf transceivers realize this functionality using Minimum-Shift Keying (MSK) with frequency modulation [58]. This puts such transceivers in the same boat as most BLE transceivers that similarly use binary frequency-shift keying (BFSK)⁴. Thus, it would be accurate to say that CT-based protocols are dominantly tested for BFSK based wireless technologies, something that also defines the scope of this paper. For a comprehensive overview of BFSK-based CT protocols, please refer to [10, 36, 73].

Early works such as Glossy [29] showed that, when using frequency-based modulations, if nodes are sufficiently synchronized, then transmissions of the *same data* will sufficiently align and the packet will be correctly received with cooperative gain. Meanwhile, later works such as Chaos and Crystal have shown that transmissions of *different data* greatly benefit from the 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) *Time synchronization errors*. Concurrent transmitters are not perfectly synchronized. This introduces additional inter-symbol interference (ISI) than what would be caused by multipath fading alone. ISI happens when different symbols overlap on the air on the receiver side. To minimize this effect, packet transmissions must be triggered within a time interval ideally lower than half the symbol period [24]. Even when transmitters are perfectly synchronized, ISI would still be present due to different time-of-flight from each transmitter to a receiver. Nevertheless, this typically adds only a small fraction of the symbol period for short-range wireless personal area network technologies such as Bluetooth or IEEE 802.15.4, and is thus not a dominant factor.
- (2) *Frequency synchronization errors*. Transmitters do not only send data over the air at slightly different times, but they also experience a mismatch between the frequencies of their carrier signals due to their imperfect crystal oscillators. When these carrier signals overlap on the air at the receiver side, the resulting waveform has alternating periods of high and low amplitude signals, giving rise to constructive and destructive

³The capture effect allows a receiver to correctly receive a packet despite multiple (and with potentially *different* data) CT when the incoming RF signals satisfy certain power and timing constraints. The latter relate to the power and time with which the stronger signal is received (compared to all other signals), and are linked to the concepts of capture threshold and window [94]. With CT, the (power) capture threshold is greatly extended compared to the classical capture effect by synchronously sending *the same* data. [89].

⁴For the same reason, it is very common today to see both standards to coexist in a single RF integrated circuit.




Fig. 2. **Beating manifests in many forms depending on the number of concurrent devices, their relative frequencies, and their channel gains.** These simplistic simulations are intended to demonstrate the occurrence of beating when multiple signals are superimposed. Sine waves are generated at 15 Hz steps, between 400-500 Hz, and are superimposed. If there is no dominant signal (i.e., the signal powers are equal), then *strong* beating attenuates a signal toward zero in a beating valley. Conversely, the presence of dominant signal (simulated as 5x the amplitude of the other signals) *weakens* the beating. As the number of transmissions increases, the resultant amplitude modulation gets more complex, with a tendency to narrower periods of varying amplitude.

interference (beating). This is called the *beating effect*. As shown in Fig. 1, with *two* concurrent transmitters, the waveform's envelope has a sinusoidal shape with periods of constructive and destructive interference, whereas it features more complex shapes when *more than two* transmissions overlap [53]. While potentially introducing a certain degree of energy gain during peaks, beating increases the chance of bit errors during low-energy periods (valleys). The beating effect is a major focus of this paper, as it has not been explored in sufficient depth in earlier literature.

Thus, intuitively, the beating effect under concurrent transmissions has the potential to cause significant packet losses. The extent of this partly depends on two factors:

- *How destructive is the interference?* When concurrently transmitted signals superimpose out of phase at the receiver, these tend to cancel each other out. If the envelope of the signal drops below the sensitivity threshold of the receiver, albeit for a short duration, the receiver sometimes cannot decode symbols correctly.

We refer to this as *strong beating*, in contrast to *weak beating* in which case the envelope stays above the sensitivity threshold. Whether beating will manifest as strong or weak depends on the relative difference in received signal energy. When signals with similar power superimpose out of phase, strong beating is most likely to happen. A high power difference would typically result in a weak beating or capture effect.

- *How wide is the destructive interference?* When concurrently transmitted signals superimpose out of phase for a duration that is long in relation to the packet on-air time, we refer to *wide beating* (as opposed to *narrow beating*). Beating will be randomly narrow or wide depending on the relative carrier frequency offset (CFO) between transmitters. Smaller offsets result in wide beating; larger offsets, instead, result in narrow beating.

That being said, it is unlikely that a radio will experience continuous (i.e., long periods of) destructive interference. Therefore, it should be noted that if the transmission lasts long enough, something will eventually be received. Following this logic, if you cannot accurately synchronize, or do not wish to synchronize, then beating potentially avoids this interference when the carrier frequencies of senders are (almost) equal and the phase offset is π . Early work on CT focused on OOK modulated physical layers [63]. In that case, the bit time is longer than the beating period, so one would eventually see a beating ‘valley’ during one bit. However, the move to IEEE 802.15.4 and BLE 5 allows the beating to avoid continuous destructive interference, as the bit time is (usually) shorter than the beating period.

In Fig. 2a we categorize beating as *wide and strong*, *wide and weak*, *narrow and strong*, or *narrow and weak*. However, based on this classification, it is not straightforward to provide qualitative statements on packet error rates. It is useful here to highlight several other factors that dictate the reliability of CT-based protocols in presence of beating:

- *Error correction.* This refers to the robustness of the coding scheme employed by the underlying PHY, i.e., to the ability of such coding scheme to cope with the beating pattern.
- *Packet length.* Reliability generally reduces as packet length increases. But an unpredictable beating length affects the error rate for narrow and wide beating differently. If the beating period is very narrow then several peaks and valleys can appear within a packet transmission and if it is very wide then a packet can be completely shadowed within a valley.
- *Number of concurrent transmitters.* Receivers listening from more than two concurrent transmitters experience more complex beating patterns as shown in (Fig. 2b to Fig. 2e). The period between valleys decreases, and peaks and valleys of different amplitudes appear. Remarkably, stronger peaks and deeper valleys may appear, potentially disrupting the reception even when there is a dominant transmitter. If the combined strength of all the other, non-dominant, transmissions is comparable to the strength of the dominant one, which is typical in dense topologies, the reception may still be heavily affected.

With background into CT-based protocols and factors affecting their performance, it is clear that reliability of a *single CT link* depends on not one but multiple factors. We next provide insights into *link-level* performance through simulation in Sect. 3 and then use empirical experiments in Sect. 4. Moving from *link* to a *network wide* performance, we resort to large-scale testbed experiments using known CT protocols in Sect. 5.

3 IMPACT OF PHYSICAL LAYER ON CT PERFORMANCE: ANALYSIS AND SIMULATION

While synchronization errors can be greatly reduced by a proper low-level design of 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) – a random deviation from the nominal transmission frequency that occurs due to the tolerance of the quartz crystals used in the radio oscillator, usually expressed in parts per million (ppm). For comparison, in classical MIMO communication systems CT

come from the same local oscillator, so multipath effects dominate instead of beating, which is a unique property of systems using independent transmitters.

To better analyze the impact that different PHYs have on the performance of CT under beating, we simulate different low-power communication systems supported by the nRF52840 platform using MATLAB to obtain the average Packet Error Rate (PER) vs. Signal-to-Noise Ratio (SNR) for Concurrent Transmissions from two nodes (a scenario later referred to as CT_2 in this paper) recovered with a non-coherent BFSK receiver, as in [25]. We assume additive white Gaussian noise (AWGN) and no synchronization errors. We repeat this for different uncoded (BLE 5 1M and 2M) and coded (BLE 5 500K, BLE 5 125K, and IEEE 802.15.4) PHYs, for different oscillator inaccuracies (which result in either wide or narrow beating), and for power deltas. Both BLE 5 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 [84]. In addition, the 125K PHY adds a Manchester pattern mapper of four elements per coded bit⁵. The IEEE 802.15.4 coding for devices operating in the 2.4 GHz ISM band is based on DSSS.

We define the Relative Frequency Offset (RFO) – which determines the beating frequency – as the difference between the CFO of each individual transmitter. We further define 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:

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

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

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

The BLE standard requires the CFO to be within ± 150 kHz [10], which results in RFOs lower than 300 kHz. Therefore, the RFO is always lower than the symbol frequency (i.e., 2 MHz in BLE 5 2M and 1 MHz for the other three BLE 5 PHYs). Similarly, the IEEE 802.15.4 standard limits the CFO to ± 100 kHz [36].

Experimental setup. We build a simulator using MATLAB and release its code open-source⁶. Specifically, we follow these steps:

- (1) We start with the baseband model, typically available in MATLAB for common communication systems, and use it to generate a random packet stream. Since we want to accurately characterize the beating effect within a symbol period, we obtain our baseband-equivalent stream, $B(t)$, with 8 samples per symbol.
- (2) Then, we add to the baseband signal the effect of the amplitude distortion that occurs when two concurrent transmissions overlap in the air sending the same bitstream. When they are both received with the same energy ($\Delta P = 0$ dB), this is equivalent to introduce a periodical amplitude modulation of period equal to twice the beating period [25]:

$$B_{CT_2}(\Delta P = 0 \text{ dB})(t) = 2B(t) \left| \cos(\pi t/T_{\text{Beating}}) \right| \quad (4)$$

If we consider energy differences between the two CT, the baseband-equivalent signal can be modeled as [25]:

$$B_{CT_2}(t) = B(t) \left(1 - A_2/A_1 + 2(A_2/A_1) \left| \cos(\pi t/T_{\text{Beating}}) \right| \right) \quad (5)$$

Where A_1 and A_2 denote the amplitude of the received signals, and $A_1 > A_2$:

$$P_{R1}/P_{R2} = A_1^2/A_2^2 \quad (6)$$

For more than two CT, the model of the amplitude distortion of the baseband envelope is much more complex, as sketched in Fig. 2b to Fig. 2e. For simplicity, we will stick to two CT throughout the simulations,

⁵When using BLE 5 125K PHY's Manchester Pattern Mapper, a '0' is translated into '0011', whereas a '1' is translated to '1100' [9, 10].

⁶<https://github.com/ADEscobar/ct-simulator.git>

which gives illustrative-enough results for typical scenarios that can be modelled as having a dominant transmission.

- (3) Finally, we recover the original bitstream by demodulating the distorted baseband signal, $B_{CT_2}(t)$ with a suitable receiver. For the particular case of a simple non-coherent BFSK receiver, the bit error rate (BER) for $\Delta P = 0$ dB can be obtained analytically [25]:

$$BER_{CT_2\text{ BFSK}}(\Delta P = 0\text{ dB}) = \frac{1}{2} \exp(-E_{b0}/N_0) I_0(-E_{b0}/N_0) \quad (7)$$

where $I_n(z)$ is the modified Bessel function of the first kind. Analytical derivations of the BER when introducing energy imbalances, or for more complex communication systems such as the coded modes of BLE 5 and IEEE 802.15.4, are very complex. Obtaining the analytical packet error rate (PER) is also difficult even for the most simple case (BFSK), since bit errors are not randomly distributed and tend to appear during the energy valleys. Similarly, the shape of the envelope becomes too complex to be analytically tractable for more than two CT.

Fig. 3 and Fig. 4 present the results of our simulation. In general, beating increases the error rate in typical low-noise conditions; only in very harsh environments may reception chances improve due to the path diversity and energy gains during the peaks. Based on these results, we derive the following observations:

- (1) **Impact of beating.** We first compare the results obtained with concurrent transmissions from 2 nodes (CT_2) 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), CT_2 experience in some cases 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. 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 packet⁷, 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 (IEEE 802.15.4, BLE 5 125K and 500K), since the former benefit from faster transmissions. With wide energy valleys burst errors are frequent. The convolutional coding used by BLE 5 125K and 500K, and the DSSS used by IEEE 802.15.4 are effective against discretely distributed one-bit errors, but not to correct contiguous errors. As a result, BLE 5 125K is the worst performing PHY, since it results in the longest packet duration.
- (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. BLE 5 1M and 2M greatly struggle and feature a very high PER, since errors are frequent during the unavoidable narrow valleys, and since uncoded PHYs do not feature any built-in correction mechanisms. Contrarily, coding modes are very effective in narrow-beating conditions, since narrower valleys introduce shorter (potentially one-bit) error bursts that can be effectively corrected. BLE 5 125K is the best performing PHY, since it provides the most robust error-correction mechanism, while BLE 5 500K and IEEE 802.15.4 perform similarly.
- (4) **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, as can be seen in the simulations already with only 1 dB difference.

⁷ T_{Packet} depends on the PHY and it is computed as $T_{Packet} = B \cdot (1/DR)$, with DR being the effective data rate of the particular PHY and B being the length of the packet in bits.







(a) Wide & strong beating (CT_2, RFO=500Hz, $\Delta P=0$ dB)(b) Wide & weak beating (CT_2, RFO=500Hz, $\Delta P=1$ dB)(c) Narrow & strong beating (CT_2, RFO=10kHz, $\Delta P=0$ dB)(d) Narrow & weak beating (CT_2, RFO=10kHz, $\Delta P=1$ dB)(e) Wide beat. & capture effect (CT_2, RFO=500Hz, $\Delta P=6$ dB)(f) Narrow beat. & capture effect (CT_2, RFO=10kHz, $\Delta P=6$ dB)

Fig. 3. **Simulation results showing the impact of beating on the PER.** Simulating two concurrent transmitters with 30 byte packets. PER of a single transmitter (no beating) is also displayed for comparison. Differences in beating periods (wide or narrow) and power deltas have a significant impact on how beating affects performance. Different node pairs (with different RFOs) may therefore experience very different PERs even when operating with the same PHY and under similar conditions.




Fig. 4. **Simulation results showing the impact of beating frequency (RFO) on the PER.** Simulating two concurrent transmitters with 30 byte packets. Uncoded PHYs (BLE 5 1M and 2M) perform better under low beating frequencies, while coded modes work optimally under high beating frequencies. Coded modes (IEEE 802.15.4, BLE 5 125K and 500K) experience an error peak in the middle of the frequency range, after which the coding becomes effective and the error decreases rapidly. Already with 1dB power delta (weak beating) the PER decreases greatly and it is constantly below 0.01 for the coded modes

- (5) **Capture effect ($\Delta P > 6\text{ dB}$), Fig. 3e and 3f.** For greater power dissimilarities, the well-known capture effect [31], i.e., the ability of the receiver to recover the message of the dominant transmission, typical of frequency modulations, kicks in. The PER tends to resemble the no-beating scenario as the power delta increases.
- (6) **Impact of the beating frequency, Fig. 4a and 4b.** Coded modes perform optimally for high beating frequencies (i.e., larger than 10 kHz), since narrow valleys tend to create single-bit errors that can be properly corrected. In the low range, energy valleys span several symbols, errors cannot be corrected, and the uncoded modes (with shorter symbol durations) perform better. For weak beating, coded modes perform generally better for every beating frequency at the considered SNR of 25 dB.

In real-world networks, all six 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 should have robust behavior in all scenarios, since controlling operating conditions within a wireless network is challenging. As presented later with the experimental results, links can randomly have very different PERs depending on the strength and frequency of the beating they experiment; and the optimal PHY selection depends on the particular conditions of that link, like node layout, external noise level, and pairwise RFOs.

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. Uncoded PHYs are optimal in wide-beating conditions, since shorter packet durations combined with packet repetitions are likely to randomly trigger a successful transmission spanning an energy peak. Conversely, in narrow-beating conditions, coded PHYs behave optimally, since they are able to correct sparse errors. BLE 5 1M and 2M are optimal for wide beating, but they behave poorly under narrow beating, while the opposite holds true for BLE 5 125K. BLE 5 500K and IEEE 802.15.4 constitute an interesting trade-off between data rate and PER

in both narrow and wide beating. A system optimally designed to operate under beating should feature burst error-correcting mechanisms, such as interleaved coding, which is missing in the analyzed low-power protocols.

To focus the discussion on the beating itself, different noise patterns, such as intermittent jamming or additional source of errors, such as synchronization mismatches between the CT, have not been considered in our simulation, but they have an impact in the testbed results presented later, particularly if the retransmissions are not synchronized within half the symbol period [24], e.g., $0.5 \mu\text{s}$ for BLE 5 1M.

Next, we confirm our simulation results on real hardware using *over-the-air* experiments. Based on the simulations, there is the potential to experimentally identify if a link is operating under narrow or wide beating by comparing the PER performance in the different radio modes. Since the initial phase of the beating is random for each packet transmission, we expect the errors to be randomly scattered within a packet transmission. Surprisingly, experimental results show differently, and we experience a fluctuating error probability when we map the distribution of bit errors across a packet with a frequency matching the expected beating pattern, which we think is coming from the calibration that the receiver performs during the reception of the packet preamble. This is a simple and effective way to obtain and confirm the impact of the beating frequency without requiring additional hardware.

4 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) [29]. More recent works [24, 45], in addition to this paper's analysis in Sect. 3, 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. 3. 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 when connecting a small number of CT devices to an SDR using coaxial cables [53]. Although these efforts help to better explain some of the processes underpinning CT-based 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. 3), 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. Furthermore we extend the D-Cube testbed with TempLab temperature control modules [12], and observe how changes in temperature at transmitters and receivers can affect the impact of beating. We perform all experiments using the *Atomic-SDN* CT stack developed for the EWSN 2019 Dependability Competition [3, 5, 6].

4.1 Experimental Setup

We configure the D-Cube testbed in order to provide a single-hop scenario for up to 12 concurrently-transmitting nodes and a fixed receiving node. Fig. 5 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 .

Over-the-air experiments. We firstly perform over-the-air experiments to compare the effect of beating on different PHYs, as well as the number (density) of concurrent transmitters. Node I is configured to periodically generate and transmit a pseudo-random payload every 250 ms, which is logged before each transmission in T_{slot_1} , alongside an 8-byte CT header. When using IEEE 802.15.4, this corresponds to a 119 B payload (due to the 127 B maximum transmission unit limitation), while a payload of 200 B is used when comparing beating across the various BLE 5 PHYs, which allows sufficient time to capture beating effects while reasonably limiting transmission time when using the BLE 5 125K PHY. At the receiving node R , the byte arrays of all correctly and incorrectly received packets are logged in T_{slot_2} . A direct comparison of transmitted and received byte arrays subsequently allows observation of the following metrics.

- (1) *Bit error distribution* – error distribution across a packet exposes 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)* – an energy minimum during a preamble results in the radio discarding the packet.

Collection of the bit error distribution allows passive observation of beating-induced errors, and is the main focus of this section. Additionally, while PRR or PER are commonly used metrics for measuring protocol performance (typically, $PRR = 1 - PER$) it is useful to also consider PLR when comparing different physical layers. Specifically, periods of beating energy minima can potentially result in receivers unable to detect a packet preamble i.e., $PRR = 1 - (PER + PLR)$.

We run this setup across the IEEE 802.15.4 and all the BLE 5 PHYs for concurrent transmissions from 2, 3, 4, 6, 8, 10, and 12 nodes at -8 dBm⁸, where we define the CT density as CT_2, CT_3, ... CT_12 respectively. Each experiment runs for $\approx 18K$ transmitted packets, representing over 100 total hours of experimental data.

Temperature-controlled experiments. Furthermore, we perform experiments on temperature-controlled nodes, connected via coaxial cabling, to examine the impact of temperature on the beating effect (and, consequently, on CT performance). We utilize the TempLab design outlined in [12], which are equipped with an extended variant of D-Cube's observer modules – with additional input for a DS18B20+ temperature sensor as well a controllable fan. A Raspberry Pi 4 Model B implements a closed-loop control system with an individual PID controller per node directly actuating an incandescent heat lamps operated on a wireless dimmer circuit mounted directly above each node. The dimmer uses Z-Wave technology, which is based on a sub-GHz radio, and thus does not affect the communication of the nodes. Fig. 6 shows a single observer module enhanced with TempLab, where the porous concrete isolation retains little heat and serves as fire retardant. Each observer module reports its temperature measurements from the DS18B20+ temperature sensor to a tempLab topic on the testbed's message broker. The Raspberry Pi 4 Model B running the control system subscribes to these measurements from the message broker and feeds them into the PID controller responsible for the adjustment of the corresponding observer modules' temperature. Using the temperature profile as set value, the PID either turns on the heating lamp using the dimmer in order to increase the temperature, or sends back commands through the broker to the observer module so it actuates the attached fan in order to decrease the nodes' temperature. We subsequently use this setup to measure the impact of temperature on beating patterns between node pairs: we do so by varying the temperature at a *single* transmitter between 30°C and 75°C at 1°C intervals. We transmit over BLE 5 500K as the beating

⁸We select a transmission power of -8 dBm, as it allows to capture sufficient bit errors within a reasonably short time window, while minimizing experimentation time.




Fig. 5. **Experimental setup.** An initiating node I synchronizes transmitting nodes, allowing a receiving node R to examine the 1-hop CT (CT₂ – CT₁₂) transmissions.





Fig. 6. **Enhanced observer module in D-Cube.** An nRF52840DK node equipped with TempLab's temperature sensor, heating lamp, and cooling fan. Each observer module is controlled by an individual PID instance on a central control system.

sensitivity of this PHY provides clearer beating patterns within the histogram (as we show in Fig. 10) and is thus more quantifiable by an FFT. Each 1°C interval lasts 10 minutes, where we send 255 byte packets every 333ms⁹, for a total of ≈1800 samples per interval. This experiment is run over all possible transmitting pair / destination combinations between four nodes {1, 2, 3, 4}, with each transmitting pair connected to the a single receiving node using coaxial cables to remove multipath effects, and we make these data publicly available¹⁰.

4.2 Results

Beating effect on different PHYs. Fig. 7 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

⁹In contrast to experiments comparing beating over different PHYs, using *only* BLE 5 500K ensures that use of the full MTU doesn't take excessive time, and allows collection of a greater number of samples.

¹⁰<https://zenodo.org/record/5882965>




Fig. 7. A 2 kHz beating frequency can be seen from the *same* CT pair across multiple *different* PHY layers. A histogram of bit errors across many packets shows the beating frequency. Low data-rate PHYs experience *narrower* beating as they span (in time) a greater number of error peaks. The coding used in BLE 5 125K handles these errors.




Fig. 8. PRR, PER, and PLR as CT density increases from a single transmitter (No_CT) to concurrent transmissions from twelve nodes (CT_12). Different PHYs handle the *complex* beating frequencies (as described in Sect. 2) with varying degrees of success.

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* CT_2 pair is ≈ 0.5 ms (a frequency of 2 kHz).

Fig. 9 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. 3, the higher data rate PHYs (BLE 5 2M and 1M) 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




Fig. 9. PRR and PER for the subplots shown in Fig. 7. The CT_2 pair from this figure experiences *significantly strong* beating over a 200 byte packet.




Fig. 10. Individual CT_2 pairs experience *different* beating frequencies depending on their RFOs. This can be seen from examination of the bit error distribution across nine *different* CT_2 pairs when using the BLE 5 500K PHY.

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. 9). 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. 8. Experiments were run across all PHYs for a single transmitter (*no CT*), as well as increasing CT density from CT_2 to CT_12. Plots represent an average of multiple experiments run with randomly selected CT forwarders per experiment, while the same




Fig. 11. **Wide and strong beating frequencies result in considerable losses.** CT reliability (PRR and PER) for the nine CT_2 pairs shown in Fig. 10 shows the coded BLE 5 500K PHY struggles when experiencing *wide and strong* beating.

pseudo-random forwarding set remains consistent across each PHY. This averaging eliminates bias due to *narrow and strong* beating experienced by CT_2 pairs such as Figs. 10e and 10h. Reliability drops at CT_3 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 [53] and with our analysis in Sect. 3. While the IEEE 802.15.4 PHY still performs well at CT_3, its PRR also drops significantly at CT_4. 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. 10 examines the bit error distribution for nine different CT pairs on the BLE 5 500K PHY – chosen due to its beating sensitivity, which allows clear beating patterns within the figures. 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. 10a and 10b, 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. 11 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.

Calculating RFO between devices. Figs. 12a–d present histograms of the bit errors generated by four independent CT_2 device pairs transmitting over BLE 5 500K. A clear sinusoidal beating pattern emerges on each of the four pairs, with Fig. 12a exhibiting *wide and strong* beating, while Fig. 12d suffers from *narrow and strong* beating (as categorized in Sect. 2). Interestingly, by sampling the signal and applying an FFT across the bit error data, it is possible to calculate this beating frequency generated as the result of the RFO between the devices. Figs. 12e–h show the FFT and the calculated RFO for each of the four CT_2 pairs depicted in Fig. 12, where we calculate a low frequency of 496Hz for Pair1, while Pair4 exhibits a relatively high beating frequency of 7946Hz. Note that, while this process has been performed statically for the purposes of this paper, the computing resources available in modern IoT chipsets (such as the nRF52840, or the dual-core nRF5340) introduce the possibility of performing such calculations on-the-fly during runtime.

Complex beating resulting from multiple RFOs. Sect. 2 introduced the concept of *complex* beating when more than two devices are simultaneously transmitting. Fig. 13a shows the complex beating pattern generated by three transmitting devices (CT_3), as observed through the histogram of bit errors at the receiver. Fig. 13b shows the FFT of the treated signal, where the two highest peaks indicate the frequency of two waveforms at 248 Hz and 3720 Hz and correspond to the multiple RFOs between the devices. We observe that *these two distinct*




Fig. 12. The RFO between devices can be quantified with an FFT. The beating frequencies of four *different* CT_2 pairs transmitting over the BLE 5 500K PHY are resolved by applying an FFT over the beating histogram.




Fig. 13. **RFOs can be discerned even in complex beating patterns.** Complex beating observed at a receiver as the result of oscillator offsets between three transmitter nodes (CT_3). Multiple peaks in the FFT correspond to the frequencies observed in the histogram.

sinusoidal beating patterns are clearly evident as opposed to a single peak present in the CT_2 FFTs in Fig. 12. Analysis of these complex beating waveforms could therefore potentially indicate the number of concurrently transmitting nodes.

Effect of temperature on beating. Since common off-the-shelf transceivers use non temperature compensated crystal oscillators, as temperature increases, the frequency is expected to decrease [93]. Since the CFO of a device is linked to this oscillator frequency, a change in temperature will result in a CFO change on a transmitting node. This subsequently affects the RFO between two concurrently transmitting nodes, and therefore the perceived




Fig. 14. **Effect of temperature on beating frequency for node pairs in Fig. 16.** CFO changes induced by temperature are shown to affect beating frequency, regardless of the receiver.

beating pattern at the receiver. For example, if we have two concurrently transmitting nodes $TX1$ and $TX2$, where the carrier frequency of $TX1$ is greater than that of $TX2$ (i.e., $CF_{TX1} > CF_{TX2}$), increasing the temperature at $TX1$ will *decrease* its oscillator (carrier) frequency and therefore *reduce the RFO between the two devices*. As the temperature at $TX1$ continues to increase, the CF_{TX1} may eventually drop below that of CF_{TX2} (i.e., $CF_{TX1} < CF_{TX2}$). At such a point the RFO will no longer decrease, but will instead increase. Figs. 14, 15, 16, and 17 examine the effect of temperature on pairs of concurrently transmitting nodes, where the temperature of one




Fig. 15. **Effect of temperature on Packet Error Rate (PER) for node pairs in Fig. 16.** Beating frequency changes induced by temperature changes can be seen to directly affect reliability at the receiver. While PER is also governed by other factors such as noise floor, PER changes in response to temperature generally follow similar curve regardless of the receiver.

node is gradually increased (between 30°C to 75°C at 1°C intervals¹¹). This is done for all possible permutations of two transmitting nodes to a single receiving node within in the set of four nodes {1, 2, 3, 4}. In Figs. 14–17, $X_T^H + Y_T - Z_R$ indicates a transmitting node pair $X_T + Y_T$ (where the subscript ‘T’ denotes transmission and

¹¹As TempLab does not offer cooling capabilities allowing testing below 30°C, we select a temperature range emulating the variations in both ambient and direct sunlight scenarios during summer. Note that what is relevant is the *difference* in temperature and not the absolute temperature, i.e., the CFO changes induced by temperature are expected to also affect beating frequency when temperature goes below 30°C.




Fig. 16. Packet error rate vs. temperature induced beating frequency for all transmitting pair / destination configurations between a set of four nodes $\{1, 2, 3, 4\}$. The PER response w.r.t. frequency is shown to be within the bounds expected by simulations depicting a 14 dB E_b/N_0 channel where transmission power delta is between 2 dB to 4 dB.

node X is heated and denoted by the superscript ' H ') and the CT are recorded at the receiving node Z (where the subscript ' R ' denotes reception). For these particular experiments, coaxial cabling was used to reduce the multipath effects that could potentially distort the beating histogram.

In Fig. 14 we observe the beating frequency curves follow similar trends as a function of temperature. In each case, we heat one of the T_x nodes while keeping the other fixed and send to two different destinations (i.e., the remaining two nodes in the set $\{1, 2, 3, 4\}$). We then repeat the experiment while heating the other T_x node. As the crystal oscillator of each node has unique imperfections of the manufacturing process, each response is slightly different. Figs. 14a to 14f show each of the 6 possible transmitting node pairs in the set of nodes $\{1, 2, 3, 4\}$. We further observe that the response of the beating frequency remains consistent regardless of the destination, and conclude that the beating frequency at a receiving node depends *solely* on the CFO of the transmitting nodes.




Fig. 17. **Beating frequency vs. Packet Error Rate (PER) for all node pairs in Fig. 16.** The black curves (solid, dashed, and dotted) indicate the $E_bN_0=14$ dB simulation results for $\Delta P=\{0,2,3,4\}$ dB, while the solid black vertical lines indicate the approx. beating frequency range detectable by our technique – running an FFT over the bit error histogram.

Fig. 15 shows the Packet Error Rate (PER) at each receiver during those experiments. It can be seen that, as with beating frequency in Fig. 14, the PER response follows a curve w.r.t. temperature. However, while the rate of change of this response is agnostic to the receiving node (i.e., the shape of the curve remains consistent, regardless of which node receives), Figs. 15c and 15e, in particular, indicate that other factors such as the noise floor at the receiving node may introduce a constant offset in the error rate. In these cases, we see considerable difference in the PER (around 60%) when transmitting to different receivers. Interestingly, Figs. 14c and 14d exhibit large erroneous spikes at around 37-62 degrees in Fig. 14c and 40-56 degrees in Fig. 14d. These can be

explained when comparing to their respective plots in Fig. 15. In Fig. 15c there is a high PER for the $1_T^H + 4_T$ node pair (regardless of receiving node), where a high noise floor has reduced the definition of the waveform in the histogram, reducing the ability of the FFT to quantify the beating frequency. In Fig. 15d node pair $2_T^H + 3_T$ (again regardless of receiving node) experiences relatively high reliability, reducing the number of samples available for the FFT.

Fig. 16 compares the results from Fig. 14 and 15 to show the PER response w.r.t. frequency changes (induced by temperature), while Fig. 17 breaks down this comparison into the node pairs used in those previous figures. We use the simulator from Sect. 3 to plot the expected PER response to beating frequency for various power delta (ΔP) and normalised SNR ($E_b N_0$) values. Specifically, we simulate an SNR of 25 dB, 18 dB, and 14 dB as a reference for high to low SNR values. Furthermore, while Nordic nRF52840 radios may deviate from their set transmit power by ± 4 dB across a Gaussian distribution (meaning that any two transmitting nodes may potentially have a power delta of up to 8 dB), under normal conditions we assume this will generally be ± 2 dB [55]. We observe that the trend when plotting frequency against the PER generally matches that of the simulations. In particular, the curves most accurately follow the 14 dB $E_b N_0$ simulations between 2 dB to 4 dB power delta. We therefore include these as reference in the breakdown of results into node pairs in Fig. 17. Furthermore the vertical black lines on the figures indicate the approximate granularity of our beating estimation characterization (generating a histogram of bit errors and running an FFT over the resulting waveform) – i.e., our technique is able to provide adequate resolution to discern beating frequencies approximately < 250 Hz and > 20000 Hz in BLE 5 500K over 250 bytes packets, which results in 2000 data points sampled at 500 kHz.

4.3 Key Observations

Recent literature has theorized that the beating effect should have a significant impact on CT performance [53, 66]. 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. 10, 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 [10]. 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. 7, 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. 9 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. 3). 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.

RFO between devices can be quantified by observation of the beating waveforms. Fig. 12 shows how, by applying an FFT over the bit error histogram observed at a receiver, it is possible to calculate the beating frequency, and that this frequency is unique to each CT pair. Fig. 14 subsequently shows that the frequency curve w.r.t. temperature is similar regardless of the receiving (observing) device. This has important security and privacy implications. Given the increasing performance capabilities of modern low-power IoT radio chipsets, this technique could potentially be used to identify individual devices, either for authentication, or as an attack vector.

Complex beating patterns indicate the number of transmitters. Fig. 13 demonstrates that it is possible to apply an FFT even over complex beating waveforms, and discern the *number* of devices within a transmission. Not only does this again have important security implications in that nodes are no longer ‘hidden’ within a flood, but it could potentially be used to improve CT performance. Specifically, as per Fig. 8, it is notable that the performance of CT can potentially be severely affected by a high degree of nodes within a single transmission. As such, knowledge of the number of participating nodes could allow the design of new protocols that reduce the number of active nodes within a local area, depending on the underlying PHY, thereby improving performance.

Beating frequencies are highly sensitive to temperature. The sensitivity of crystal oscillators frequencies is well documented [93]. As beating frequencies are ultimately linked to the relative carrier offset between devices, temperature variations also affect CT-induced beating frequencies. Fig. 14 and Fig. 15 explore how temperature changes at individual transmitters can increase or decrease the beating frequency, and subsequently the overall PER of concurrent transmissions. Importantly, Figs. 16 and 17 ultimately demonstrate that this impact is captured within our simulation and modeling from Sect. 3, and that the temperature sensitivity is predictable within known channel bounds. While chipsets such as the nRF52840 don’t permit fine-grained tuning of the carrier frequency, alternative platforms (such as an SDR) could potentially correct the CFO to improve error rates.

5 CT PERFORMANCE OVER DIFFERENT PHYs:

EXPERIMENTAL EVALUATION WITH RF INTERFERENCE

This section presents an experimental study about the impact of different PHYs on the performance of CT-based protocols in the presence of RF interference. Specifically, we evaluate three CT-based data dissemination protocols (i.e., Glossy [29], Robust Flooding (RoF) [48], and RoF operating on a single channel (RoF (SC))), as well as three CT-based data collection protocols (i.e., Crystal [37], multi-channel Crystal (Crystal^{CH}) [38], and multi-channel Crystal with noise detection (Crystal^{CH}_{ND}) [38]) on a large multi-hop network. The operations of these CT-based protocols are depicted in Fig. 18; we refer the reader to [94] for a detailed survey on CT-based protocols.

The original Glossy [29] 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 [48, 51]. Using this technique, it is difficult to resume a CT flood if the latter is interrupted by RF interference. An alternative approach was taken by the authors of [48, 62], 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 a 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 operating on a single channel (RoF (SC)) to observe how this protocol performs with respect to the single-channel environment used by Glossy.

Protocol type	Protocol	TX_N	Duration (ms)	Period (ms)	Epoch (ms)	Channel	Trigger mode
Dissemination	Glossy	6	100	200	-	2.480 GHz	Rx-Tx
	RoF (SC)	6	100	200	-	2.480 GHz	Rx-Tx-Tx
	RoF	6	100	200	-	2.402, 2.426, 2.480 GHz	Rx-Tx-Tx
Collection	Crystal	6	20, 45, 15	30, 57, 27	1000	2.480 GHz	Rx-Tx
	Crystal ^{CH}	6	20, 45, 15	30, 57, 27	1000	2.402, 2.426, 2.480 GHz	Rx-Tx
	Crystal ^{CH} _{ND}	6	20, 45, 15	30, 57, 27	1000	2.402, 2.426, 2.480 GHz	Rx-Tx

Table 1. Summary of protocol parameters used to evaluate CT-based flooding performance on D-Cube. Collection protocols based on Crystal employ three distinct *sync* (S), *transmit* (T), and *acknowledge* (A) flooding phases with permitted maximum durations of 20, 45, and 15ms respectively, repeating every 1s (i.e., an epoch). For each phase, the periods (i.e., period of time in which flooding is allowed) are set to 30, 57, and 27 ms respectively, as depicted in Fig. 18. Dissemination protocols employ a single flooding phase with a maximum duration of 100ms, periodically repeating every 200ms.

By using Glossy as a fundamental primitive, one can easily build many-to-one data collection protocols. Crystal [37] is an example of such CT-based data collection protocol. In Crystal, the sink node initiates a synchronization flood to all the nodes in the network. Followed by the synchronization period, there are TA pairs (one transmission period and one acknowledgment period) as shown in Fig.18 (bottom). A node can flood its message in a T period. Relying on the capture effect, one of the messages can be probably received by the sink node successfully when multiple nodes have messages to flood in a T period. In the following A period, the sink node floods its acknowledgment to indicate whose message has been just received. This way, in the next T period, the received nodes would work as a relay node. The synchronization period and TA pairs make up an epoch. If there is no nodes sending messages in consecutive T periods (i.e., silent T periods), all the nodes would turn off the transceivers and no communication happens until the next epoch, thereby saving energy. In order to work under interference, Crystal^{CH} and Crystal^{CH}_{ND} [38] enable Crystal to work on multiple channels. Rather than the slot-by-slot channel hopping mechanism utilized in RoF, where nodes rapidly hop frequency *during* the flood, these protocols simply hop to a known frequency at the start of each flooding period, maintaining the same channel throughout the flood. Table 1 summarizes the protocol parameters used. As mentioned before, a node would speculate that there is no new message if there are consecutive silent T periods. However, it is difficult for a node to know whether the silent T period is “real” or it is caused by interference. Therefore, Crystal^{CH}_{ND} introduces the noise detection (ND) mechanism, where nodes detect the interference level by using RSS values if there is nothing received in current CT slot, and enter sleep mode once harsh interference is detected (i.e., the RSS values are greater than a specific threshold).

5.1 Experimental Setup

We use the D-Cube testbed [68] and all its 50 nodes¹² to evaluate in hardware three key metrics: *end-to-end reliability*, *latency*, and *total energy consumption*. For each of the protocols (Glossy, RoF, RoF (SC), Crystal, Crystal^{CH}, and Crystal^{CH}_{ND}), we consider three scenarios characterized by the absence or presence of interference, denoted as *no*, *mild*, and *strong* interference, respectively.

D-Cube’s controllable RF interference is generated by observer nodes (Raspberry Pi 3, directly attached to the nRF52840 boards) using JamLab-NG. *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 observer

¹²https://iti-testbed.tugraz.at/wiki/index.php/Topology_of_Nodes

No.	Data Dissemination Layout	Description
(d-i)	D-Cube nRF52840 data dissemination, layout # 1	Messages sent from a sink (node 206) to 5 destinations.
(d-ii)	D-Cube nRF52840 data dissemination, layout # 4	Messages sent from a sink (node 200) to 31 destinations.
No.	Data Collection Layout	Description
(c-i)	D-Cube nRF52840 data collection, layout # 1	Messages sent from 5 source nodes to a sink (node 202).
(c-ii)	D-Cube nRF52840 data collection, layout # 3	Messages sent from 47 source nodes to a sink (node 227).
(c-iii)	D-Cube nRF52840 data collection, layout # 4	Messages sent from 10 source nodes to a sink (node 219).

Table 2. **Summary of layouts used to evaluate the performance of CT-based protocols in D-Cube.** The layout numbers refer to the official scenarios supported by the D-Cube benchmarking infrastructure, which involve nodes located in regions with both low and high node densities.

node chooses a different channel, generating interference of 200 mW for ≈ 8 ms every 13 ms. Further details about the configuration of the generated interference are available in the original JamLab-NG paper [70].

Each dissemination protocol is run on D-Cube's data dissemination scenarios, i.e., those sending data from a single source to multiple destinations over a multi-hop network. Each data collection protocol is run on D-Cube's data collection scenarios, i.e., those where a sink node collects data from multiple (or even all) nodes. We further configure D-Cube to generate aperiodic messages with short (8B) and long (64B) payloads. To ensure that protocols span the minimum hop distance of the D-Cube testbed when using BLE 5.2M (known through prior experimentation), the maximum number of transmission attempts per node during a flooding period (defined as TX_N) is set to six. The duration of a flooding and the periodicity in dissemination protocols are both set to 100 and 200 ms respectively. The durations of the synchronization, T (transmission), and A (acknowledgment) periods are set to 20 ms, 45 ms and 15 ms, respectively. Furthermore, there are 12 TA pairs in each epoch, which is set to one second. In single-channel protocols (Glossy, RoF (SC), and Crystal) the radio frequency is set to 2.480 GHz (which corresponds to channel 39 in BLE and channel 26 in IEEE 802.15.4), whereas multi-channel protocols (RoF, Crystal^{CH}, and Crystal^{CH}_{ND}) hop among three frequencies (2.402, 2.426, and 2.480 GHz, which correspond to BLE's three advertising channels 37, 38, and 39, respectively). For Crystal^{CH}_{ND} we further select a noise detection of -70 dBm: this is a good choice given that the nRF52840DK exhibits RSS values of about -90 dBm when there is no activity over the air.

All the protocols (data dissemination and collection) are run on multiple D-Cube layouts to avoid making our results setup-specific. The topology of D-Cube can be regarded as two adjacent regions, one of which is a *dense* region with around 20 nodes, whilst the remaining nodes make up a *sparse* region of the network. For the data dissemination protocols, we run two D-Cube layout configurations: a single source node to (i) 5, and (ii) 31 destinations. For the collection protocols we run three D-Cube layout configurations featuring 5, 47, and 10 source nodes to a single destination, respectively. Additionally, we set the transmission power to 0 dBm for all nodes, which ensures a network diameter between 6 and 10 hops depending on the employed layout. Table 2 provides a summary of the D-Cube layout configurations used in our experiments.

All the results shown in this section utilize publicly available implementations of Glossy¹³, RoF¹⁴, Crystal^{CH15}_{ND}, which we subsequently ported to the nRF52840-DK platform supported by D-Cube.

5.2 Results

Fig. 19 summarizes the results of our experiments and provides an overview of the end-to-end reliability and latency, as well as of the total energy consumption of the analysed data collection and dissemination protocols in the absence and in the presence of RF interference.

Reliability. For the data dissemination protocols, Fig. 19a shows that, under no interference, as data rate increases, the evaluated PHYs struggle to maintain reliability. RoF exhibits the highest reliability, while the packet reception rate of RoF (SC) drops at higher data rates compared to that of 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 $0.25 \mu s$ [2].

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 means that 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. 19a 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.

Data collection protocols exhibit similar trends to those of the data dissemination ones, i.e., shorter messages are delivered more reliably, and channel hopping mechanisms improve the reliability under interference. However, data collection protocols cannot sustain a comparably high packet reception rate to that of data dissemination protocols when the interference becomes very strong, as visible in Fig. 19b. The reason for this is that a message in a many-to-one network has less flooding chances than in a one-to-many network: in fact, multiple nodes are required to send their packets to the individual sink node and the throughput of the sink node is limited even if a channel hopping mechanism has been adopted. The noise detection feature improves the reliability by 15% and 62% for short and long payloads, respectively, under strong interference, which is coherent with the results presented in [38].

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. 19c and 19d support this, and we observe significant latency jumps as interference increases. Especially for data collection protocols (Fig. 19d), the latency can be much higher (i.e., as high as four seconds or more) under strong interference when only a single channel is used. 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.

¹³<https://sourceforge.net/p/contiki/projects/code/HEAD/tree/ethz.ch/glossy/>

¹⁴<https://github.com/ETHZ-TEC/robust-flooding-a03f38a>

¹⁵<https://github.com/d3s-trento/crystal-11ef31e>




Fig. 19. Performance of CT-based data collection and dissemination protocols in the presence of RF interference. These six plots show the average reliability, latency, and total energy consumption across the data dissemination and collection layouts of D-Cube listed in Table 2.

Energy efficiency. Similarly, the energy consumption of nodes 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. Under interference, nodes take a longer time to attempt to overhear a packet. Therefore, interference also causes more energy consumption. However, the underlying PHY rate is still a fundamental factor in energy

consumption as shown in Fig. 19e. For *long* payloads the lower data rates of the coded PHYs means that the radio can take considerable time to transmit, consuming far more energy. Coded PHYs offer a lower energy consumption also for data collection protocols (as shown in Fig. 19f), although the difference is less pronounced than in the dissemination protocols. Interestingly, in the presence of RF interference, the energy consumption is often lower when transmitting longer payloads – a behaviour that is opposite to that observed in dissemination scenarios. We believe that the reason behind lies in Crystal’s design: after several consecutive silent T periods (in which nothing is received), the sink node lets all nodes enter sleep mode by sending a command in the next A period. In other words, packets with shorter payloads are more likely to survive mild interference and complete a Crystal round, which causes a higher energy consumption compared to the use of longer payloads (which abort the Crystal round earlier and hence consume less energy). In the presence of strong interference, also packets with short payloads have a lower chance to be received: the trend shown in Fig. 19f is hence flatter.

We also note that the energy consumed by the nodes in Crystal_{ND}^{CH} is about 1% higher than that in Crystal^{CH} due to the noise detection mechanism. This is necessary to enable nodes to still overhear the network during silent T slots under interference. This way, the sink node is more likely to receive the packets despite the presence of RF interference.

5.3 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. 8). 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. 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. 19a 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.

More probability of collision (less probability of capture effect) when using faster data rates. Crystal (and its multi-channel versions) significantly rely on the capture effect when multiple nodes have a (*different*) message to send in a T period. Specifically, in a T period, all the nodes would initiate a flood by sending their messages actively, and which message can be flooded to the sink node depends on the capture effect. However, when high data rate PHYs with shorter preambles are adopted (Fig. 19b), the capture effect is not triggered so frequently, and reliability decreases [27].

6 RELATED WORK

We discuss next related work and highlight how the contributions presented in this paper advance the state-of-the-art.

CT on different PHYs. After the influential work by Dutta et al. [23] and Ferrari et al. [29] were published in 2010/2011, a large number of researchers has started to study CT and develop CT-based protocols [1, 4, 14, 16, 18–22, 28, 33, 35, 37, 41, 42, 52, 59, 80, 85, 90, 92, 94]. 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 [40, 49, 83], as well as sub-GHz short-range [7, 46] and long-range technologies [47, 82].

A number of works have recently focused on studying the feasibility of CT on BLE [2, 53, 64, 66]. Specifically, Al Nahas et al. [2] 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 [53]. Schaper [66] 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 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 [27, 38, 48, 50], as also highlighted in the context of the EWSN dependability competition series [11, 67, 69]. To sustain a dependable performance under interference, CT-based solutions have been enriched, among others, with mechanisms such as local opportunistic retransmissions [29, 78, 79], channel-hopping [20, 26, 38, 48, 72], network coding [21, 22, 34, 52, 91], noise detection [38, 50], stretched preambles [27], data freezing [50], as well as an improved understanding of the network state [13, 15, 65].

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 and dissemination 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 analysing 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. [29] have simulated CT signals with Matlab and explained how accurately packets should be aligned in order to design reliable protocols. Other works [61, 86, 87] 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 [45]. Specifically, in [45], 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 [53].

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*. Furthermore, we also characterize the impact of environmental effects such as temperature on beating.

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 and beating. They struggle under strong wide beating, which may cause a significant drop in reliability, but provide the best overall trade-off among all PHYs. ☛ High data rate PHYs help escaping strong and wide beating, but exhibit poor performance in the presence of external RF interference. ☛ The BLE 5 125K PHY features the longest range, but performs poorly when sending long packets under external RF interference and strong wide beating. ☛ Temperature changes at transmitting nodes can effectively <i>narrow</i> or <i>widen</i> beating frequencies and subsequently affect error rates (depending on PHY). 	<ul style="list-style-type: none"> ☛ In the presence of strong external RF interference, use IEEE 802.15.4 for shorter packets and BLE 5 500K for longer packets. They provide a good combination of resilience against beating and interference, and extended range. ☛ In absence of external RF interference and with a low network density, use BLE 5 2M (or 1M) and repetitions to exploit temporal redundancy^(*). ☛ Use BLE 5 125K only for very short packets or when the extended range is a must^(*). ☛ An FFT can be periodically employed to quantify beating frequency and predictively switch PHY in response to changing temperature conditions.

Table 3. **Summary of observations on CT performance over different PHYs and corresponding recommendations.**

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

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	↑	↗	↑	↗

Table 4. **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.

7 DISCUSSION AND FUTURE WORK

An increasingly large body of work has shown that, particularly in environments with high RF interference, CT-based protocols are a valuable tool allowing low-power mesh networks to mitigate the impact of harsh RF conditions and provide robust, low-latency communication. 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 3. With modern chipsets now supporting real-time PHY switching [55, 81], these insights can guide further research into CT protocols that take advantage

of multi-PHY capabilities *at runtime*. 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 and further investigation. Foremost, analytical and mathematical modeling would help to substantiate the observations presented in Sect. 3 and 4. 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.

As shown in this paper, and succinctly summarized in Fig. 16, RFO may change over time due to oscillator temperature sensitivity – resulting in changes to the beating frequency, and subsequently the effective error rate. By maintaining a histogram of bit errors and running an FFT over the resulting waveform, it is possible to quantify the effective RFO between devices. While we have demonstrated this technique statically, the prospect of continuous monitoring of the beating effect at the receiver is intriguing, and well within the capabilities of modern dual-MCU chipsets such as the nRF5340 [56].

Furthermore, results from Sect. 4 argue that additional work is needed to not only explore how *complex* beating patterns affect CT performance (rather than the simple sine waves observed when receiving from two transmitters), but how these waveforms relate to the transmitting devices. Such capability would have important and highly impactful security repercussions. Not only could monitoring of the RFO feasibly be employed to fingerprint devices as a Physical Unclonable Function (PUF) while accounting for fluctuations due to ambient temperature changes, but filtering and analysis of beating waveforms would allow estimation of the *number* of concurrently transmitting devices. Specifically, chipsets such as the nRF5340 [56] and nRF2833 [54] offer access to I/Q values within the Constant Tone Extension (CTE) included as part of Bluetooth 5.2, which would remove the reliance on 'failed' CRC packets to examine the bit errors while also allowing more fine-grained sampling for dealing with complex beating waveforms.

Techniques such as *interleaving*, i.e., bit shuffling, (which improves the robustness of FEC with respect to burst errors – the primary culprit of beating-induced errors as discussed in Sect. 3) and *bit voting* are missing in the analyzed physical layers, and would be an effective addition to increase the reliability of CT.

Finally, while this paper has mainly focused on the negative properties of beating, specifically packet errors due to destructive interference, Sect. 2 alluded to the fact that beating isn't *necessarily* bad. While beating valleys do indeed disrupt CT communication, there is an argument that further work is needed to explore positive beating effects. Specifically, long transmissions should avoid periods of destructive interference and something will eventually be received. Exploiting beating peaks is an interesting proposition that could significantly improve the performance of CT-based protocols.

ACKNOWLEDGMENTS

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] Beshr Al Nahas, Simon Duquennoy, and Olaf Landsiedel. 2017. Network-Wide Consensus Utilizing the Capture Effect in Low-Power Wireless Networks. In *Proceedings of the 15th International Conference on Embedded Network Sensor Systems (SenSys)* (Delft, The Netherlands). ACM, 1–14.
- [2] Beshr Al Nahas, Simon Duquennoy, and Olaf Landsiedel. 2019. Concurrent Transmissions for Multi-Hop Bluetooth 5. In *Proceedings of the 16th International Conference on Embedded Wireless Systems and Networks (EWSN)* (Beijing, China). Junction Publishing, 130–141.
- [3] Michael Baddeley. 2020. *Software Defined Networking for the Industrial Internet of Things*. Ph.D. Dissertation. Dept. of Electrical and Electronic Engineering, Univ. Bristol, UK.

- [4] Michael Baddeley, Adnan Aijaz, Usman Raza, Aleksandar Stanoev, Yichao Jin, Markus Schuß, Carlo Alberto Boano, and George Oikonomou. 2021. 6TiSCH++ with Bluetooth 5 and Concurrent transmissions. In *Proceedings of the 18th International Conference on Embedded Wireless Systems and Networks (EWSN)* (Virtual event). Junction Publishing, 25–30.
- [5] Michael Baddeley, Usman Raza, Mahesh Sooriyabandara, George Oikonomou, Reza Nejabati, and Dimitra Simeonidou. 2019. Atomic-SDN: Is Synchronous Flooding the Solution to Software-Defined Networking in IoT? *IEEE Access* 7 (May 2019), 96019–96034.
- [6] Michael Baddeley, Aleksandar Stanoev, Usman Raza, Yichao Jin, and Mahesh Sooriyabandara. 2019. 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* (Beijing, China). Junction Publishing, 298–299.
- [7] Jan Beutel, Roman Trüb, Reto Da Forno, Markus Wegmann, Tonia Gsell, Romain Jacob, Michael Keller, Felix Sutton, and Lothar Thiele. 2019. The Dual Processor Platform Architecture. In *Proceedings of the 18th International Conference on Information Processing in Sensor Networks (IPSN), demo session* (Montreal, Canada). IEEE, 335–336.
- [8] Andreas Biri, Pat Pannuto, and Prabal Dutta. 2019. TotTernary - a Wearable Platform for Social Interaction Tracking. In *Proceedings of the 18th International Conference on Information Processing in Sensor Networks (IPSN), demo session* (Montreal, Canada). IEEE, 346–347.
- [9] Bluetooth Blog. [n.d.]. Exploring Bluetooth 5 – Going the Distance. [Online] <https://www.bluetooth.com/blog/exploring-bluetooth-5-going-the-distance/> – Last accessed: 2020-05-26.
- [10] Bluetooth Working Group. 2019. Bluetooth Core Specification, Revision 5.2.
- [11] Carlo Alberto Boano, Markus Schuß, and Kay Römer. 2017. EWSN Dependability Competition: Experiences and Lessons Learned. *IEEE Internet of Things Newsletter* (March 2017).
- [12] Carlo Alberto Boano, Marco Zúñiga, James Brown, Utz Roedig, Chamath Keppitiyagama, and Kay Römer. 2014. Templab: A Testbed Infrastructure to Study the Impact of Temperature on Wireless Sensor Networks. In *Proceedings of the 13st International Symposium on Information Processing in Sensor Networks (IPSN)* (Berlin, Germany). IEEE, 95–106.
- [13] Martina Brachmann, Olaf Landsiedel, and Silvia Santini. 2016. Concurrent Transmissions for Communication Protocols in the Internet of Things. In *Proceedings of the 41st International Conference on Local Computer Networks (LCN)* (Dubai, United Arab Emirates). IEEE, 406–414.
- [14] Zhichao Cao, Xiaolong Zheng, Qiang Ma, and Xin Miao. 2021. COFlood: Concurrent Opportunistic Flooding in Asynchronous Duty Cycle Networks. In *Proceedings of the 18th International Conference on Sensing, Communication and Networking (SECON)* (Virtual Event). IEEE.
- [15] Doug Carlson, Marcus Chang, Andreas Terzis, Yin Chen, and Omprakash Gnawali. 2013. Forwarder Selection in Multi-transmitter Networks. In *Proceedings of the 18th International Conference on Distributed Computing in Sensor Systems (DCOSS)* (Cambridge, MA, USA). IEEE, 1–10.
- [16] Tengfei Chang, Thomas Watteyne, and Xavier Vilajosana Pedro Henrique Gomes. 2018. Constructive Interference in 802.15.4: A Tutorial. *IEEE Communications Surveys and Tutorials* 21, 1 (Sept. 2018), 217–237.
- [17] Haiming Chen, Zhaoliang Zhang, Li Cui, and Changcheng Huang. 2016. NoPSM: A Concurrent MAC Protocol over Low-Data-Rate Low-Power Wireless Channel Without PRR-SINR Model. *IEEE Transactions on Mobile Computing* 16, 2 (2016), 435–452.
- [18] Jagnyashini Debadarshini, Sudipta Saha, Olaf Landsiedel, and Mun Choon Chan. 2020. Start of Frame Delimiters (SFDs) for Simultaneous Intra-Group One-to-All Dissemination. In *Proceedings of the 45th International Conference on Local Computer Networks (LCN)* (Sydney, Australia). IEEE, 100–111.
- [19] Jagnyashini Debadarshini, Chandra Shekhar, and Sudipta Saha. 2020. Fine-grained Frequencies for Simultaneous Intra-Group One-to-All Dissemination. In *Proceedings of the 17th International Conference on Mobile Ad Hoc and Sensor Systems (MASS)* (Virtual Event). IEEE, 473–481.
- [20] Manjunath Doddavenkatappa and Mun Choon Chan. 2014. P³: A Practical Packet Pipeline Using Synchronous Transmissions for Wireless Sensor Networks. In *Proceedings of the 13th International Symposium on Information Processing in Sensor Networks (IPSN)* (Berlin, Germany). IEEE, 203–214.
- [21] Manjunath Doddavenkatappa, Mun Choon Chan, and Ben Leong. 2013. 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)* (Chicago, IL, USA). USENIX, 269–282.
- [22] Wan Du, Jansen Christian Liando, Huanle Zhang, and Mo Li. 2017. Pando: Fountain-Enabled Fast Data Dissemination With Constructive Interference. *IEEE/ACM Transactions on Networking* 25, 2 (April 2017), 820–833.
- [23] Prabal Dutta, Stephen Dawson-Haggerty, Yin Chen, Chieh-Jan Mike Liang, and Andreas Terzis. 2010. Design and Evaluation of a Versatile and Efficient Receiver-initiated Link Layer for Low-power Wireless. In *Proceedings of the 8th International Conference on Embedded Networked Sensor Systems (SenSys)* (Zürich, Switzerland). ACM, 1–14.
- [24] Antonio Escobar-Molero. 2019. Improving Reliability and Latency of Wireless Sensor Networks Using Concurrent Transmissions. *at-Automatisierungstechnik* 67, 1 (2019), 42–50.

- [25] Antonio Escobar-Molero. 2020. *Using Concurrent Transmissions to Improve the Reliability and Latency of Low-Power Wireless Mesh Networks*. PhD thesis. RWTH Aachen University, Aachen, Germany.
- [26] Antonio Escobar-Molero, Francisco J Cruz, Javier Garcia-Jimenez, Jirka Klaue, and Angel Corona. 2016. 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)* (Granada, Spain). IEEE, 1–4.
- [27] Antonio Escobar-Molero, Javier Garcia-Jimenez, Jirka Klaue, Fernando Moreno-Cruz, Borja Saez, Francisco J Cruz, Unai Ruiz, and Angel Corona. 2019. Competition: RedNodeBus, Stretching out the Preamble. In *Proceedings of the 16th International Conference on Embedded Wireless Systems and Networks (EWSN), competition session* (Beijing, China). Junction Publishing, 304–305.
- [28] Federico Ferrari, Marco Zimmerling, Luca Mottola, and Lothar Thiele. 2012. Low-Power Wireless Bus. In *Proceedings of the 10th International Conference on Embedded Network Sensor Systems (SenSys)* (Toronto, Canada). ACM, 1–14.
- [29] Federico Ferrari, Marco Zimmerling, Lothar Thiele, and Olga Saukh. 2011. Efficient Network Flooding and Time Synchronization with Glossy. In *Proceedings of the 10th International Conference on Information Processing in Sensor Networks (IPSN)* (Chicago, IL, USA). IEEE, 73–84.
- [30] Kai Geissdoerfer, Mikołaj Chwalisz, and Marco Zimmerling. 2019. Shepherd: A Portable Testbed for the Batteryless IoT. In *Proceedings of the 17th International Conference on Embedded Network Sensor Systems (SenSys)* (New York, NY, USA). ACM, 83–95.
- [31] Cengiz Gezer, Chiara Buratti, and Roberto Verdone. 2010. Capture Effect in IEEE 802.15.4 networks: Modelling and Experimentation. In *Proceedings of the 5th International Symposium on Wireless Pervasive Computing (ISWPC)* (Modena, Italy). IEEE, 204–209.
- [32] Davide Giovanelli and Elisabetta Farella. 2018. RSSI or Time-of-flight for Bluetooth Low Energy based Localization? An Experimental Evaluation. In *Proceedings of the 11th International Wireless and Mobile Networking Conference (WMNC)* (Prague, Czech Republic). IEEE, 32–39.
- [33] Zhitao He, Kasun Hewage, and Thiemo Voigt. 2016. Arpeggio: A Penetration Attack on Glossy Networks. In *Proceedings of the 13th International Conference on Sensing, Communication, and Networking (SECON)* (London, United Kingdom). IEEE, 1–9.
- [34] Carsten Herrmann, Fabian Mager, and Marco Zimmerling. 2018. Mixer: Efficient Many-to-All Broadcast in Dynamic Wireless Mesh Networks. In *Proceedings of the 16th International Conference on Embedded Network Sensor Systems (SenSys)* (Shenzhen, China). ACM, 145–158.
- [35] Nicolas Himmelmann, Dingwen Yuan, Lars Almon, and Matthias Hollick. 2020. Concurrent Wireless Cut-Through Forwarding: Ultra-Low Latency Multi-Hop Communication for the Internet of Things. In *Proceedings of the 16th International Conference on Distributed Computing in Sensor Systems (DCOSS)* (Marina del Rey, CA, USA). IEEE, 77–85.
- [36] IEEE 802.15.4 Working Group. 2016. *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) ed.). IEEE 802.15.4 Working Group.
- [37] Timofei Istomin, Amy Lynn Murphy, Gian Pietro Picco, and Usman Raza. 2016. Data Prediction + Synchronous Transmissions = Ultra-Low Power Wireless Sensor Networks. In *Proceedings of the 14th International Conference on Embedded Network Sensor Systems (SenSys)* (Stanford, CA, USA). ACM, 83–95.
- [38] Timofei Istomin, Matteo Trobinger, Amy Lynn Murphy, and Gian Pietro Picco. 2018. Interference-Resilient Ultra-low Power Aperiodic Data Collection. In *Proceedings of the 17th International Conference on Information Processing in Sensor Networks (IPSN)* (Porto, Portugal). ACM, 84–95.
- [39] Romain Jacob, Anna-Brit Schaper, Andreas Biri, Reto da Forno, and Lothar Thiele. 2020. 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)* (London, United Kingdom).
- [40] Benjamin Kempke, Pat Pannuto, Bradford Campbell, and Prabal Dutta. 2016. 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)* (Stanford, CA, USA). ACM, 137–149.
- [41] Michael König and Roger Wattenhofer. 2016. Maintaining Constructive Interference Using Well-Synchronized Sensor Nodes. In *Proceedings of the 12th International Conference on Distributed Computing in Sensor Systems (DCOSS)* (Washington, DC, USA). IEEE, 206–215.
- [42] Olaf Landsiedel, Federico Ferrari, and Marco Zimmerling. 2013. 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)* (Roma, Italy). ACM, 1–14.
- [43] Antonio Langiu, Carlo Alberto Boano, Markus Schuß, and Kay Römer. 2019. UpKit: An Open-Source, Portable, and Lightweight Update Framework for Constrained IoT Devices. In *Proceedings of the 39th IEEE International Conference on Distributed Computing Systems (ICDCS)* (Dallas, TX, USA). IEEE, 2101–2112.
- [44] Krijn Leentvaar and Jan H. Flint. 1976. The Capture Effect in FM Receivers. *IEEE Transactions on Communications* 24, 5 (May 1976), 531–539.

- [45] Chun-Hao Liao, Yuki Katsumata, Makoto Suzuki, and Hiroyuki Morikawa. 2016. Revisiting the So-Called Constructive Interference in Concurrent Transmission. In *Proceedings of the 41st International Conference on Local Computer Networks (LCN)* (Dubai, United Arab Emirates). IEEE, 280–288.
- [46] Chun-Hao Liao, Makoto Suzuki, and Hiroyuki Morikawa. 2016. 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* (Stanford, CA, USA). ACM, 354–355.
- [47] Chun-Hao Liao, Guibing Zhu, Daiki Kuwabara, Makoto Suzuki, and Hiroyuki Morikawa. 2017. Multi-Hop LoRa Networks Enabled by Concurrent Transmission. *IEEE Access* 5 (Sept. 2017), 21430–21446.
- [48] Roman Lim, Reto Da Forno, Felix Sutton, and Lothar Thiele. 2017. 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* (Uppsala, Sweden). Junction Publishing, 270–271.
- [49] Diego Lobba, Matteo Trobinger, Davide Vecchia, Timofei Istomin, and Gian Pietro Picco. 2020. Concurrent Transmissions for Multi-hop Communication on Ultra-wideband Radios. In *Proceedings of the 17th International Conference on Embedded Wireless Systems and Networks (EWSN)* (Lyon, France). Junction Publishing, 132–143.
- [50] Xiaoyuan Ma, Peilin Zhang, Ye Liu, Carlo Alberto Boano, Hyung-Sin Kim, Jianming Wei, and Jun Huang. 2020. Harmony: Saving Concurrent Transmissions from Harsh RF Interference. In *Proceedings of the 39th International Conference on Computer Communication (INFOCOM)* (Toronto, Canada). IEEE.
- [51] Xiaoyuan Ma, Peilin Zhang, Weisheng Tang, Xin Li, Wangji He, Fuping Zhang, Jianming Wei, and Oliver Theel. 2018. 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* (Madrid, Spain). Junction Publishing, 211–212.
- [52] Mobashir Mohammad and Mun Choon Chan. 2018. 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)* (Porto, Portugal). IEEE, 72–83.
- [53] Beshr Al Nahas, Antonio Escobar-Molero, Jirka Klaue, Simon Duquennoy, and Olaf Landsiedel. 2020. BlueFlood: Concurrent Transmissions for Multi-Hop Bluetooth 5—Modeling and Evaluation. *CORR – arXiv preprint 2002.12906* (Feb. 2020).
- [54] Nordic Semiconductors. [n.d.]. nRF52833 Product Specification, v1.5. [Online] https://infocenter.nordicsemi.com/pdf/nRF52840_PS_v1.5.pdf – Last accessed: 2022-01-14.
- [55] Nordic Semiconductors. [n.d.]. nRF52840 Product Specification, v1.1. [Online] https://infocenter.nordicsemi.com/pdf/nRF52840_PS_v1.1.pdf – Last accessed: 2020-05-26.
- [56] Nordic Semiconductors. [n.d.]. nRF5340 Product Specification, v1.1. [Online] https://infocenter.nordicsemi.com/pdf/nRF5340_PS_v1.1.pdf – Last accessed: 2020-05-26.
- [57] NXP Semiconductors. 2020. K32W061/K32W041 IEEE 802.15.4 and Bluetooth LE 5.0 Wireless Microcontroller Product Data Sheet, Rev. 1.2. [Online] <https://www.nxp.com/docs/en/data-sheet/K32W061.pdf> – Last accessed: 2021-06-30.
- [58] Subbarayan Pasupathy. 1979. Minimum Shift Keying: A Spectrally Efficient Modulation. *IEEE Communications Magazine* 17, 4 (July 1979), 14–22.
- [59] Valentin Poirot, Beshr Al Nahas, and Olaf Landsiedel. 2019. Paxos Made Wireless: Consensus in the Air. In *Proceedings of the 16th International Conference on Embedded Wireless Systems and Networks (EWSN)* (Beijing, China). Junction Publishing, 1–12.
- [60] Joseph Polastre, Robert Szewczyk, and David E. Culler. 2005. Telos: Enabling Ultra-Low Power Wireless Research. In *Proceedings of the 4th International Symposium on Information Processing in Sensor Networks (IPSN)* (Los Angeles, CA, USA). IEEE, 364–369.
- [61] Vijay S. Rao, Madhusudan Koppal, RangaRao Venkatesha Prasad, T.V. Prabhakar, Chayan Sarkar, and Ignas Niemegeers. 2016. Murphy Loves CI: Unfolding and Improving Constructive Interference in WSNs. In *Proceedings of the 35th International Conference on Computer Communications (INFOCOM)* (San Francisco, CA, USA). IEEE, 1–9.
- [62] Usman Raza, Yichao Jin, and Mahesh Sooriyabandara. 2017. 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* (Uppsala, Sweden). Junction Publishing, 278–279.
- [63] Matthias Ringwald and Kay Römer. 2005. BitMAC: a deterministic, collision-free, and robust MAC protocol for sensor networks.. In *EWSN*. 57–69.
- [64] Coen Roest. 2015. *Enabling the Chaos Networking Primitive on Bluetooth LE*. Master’s thesis. TU Delft, Delft, The Netherlands.
- [65] Chayan Sarkar, R. Venkatesha Prasad, Raj Thilak Rajan, and Koen Langendoen. 2016. Sleeping Beauty: Efficient Communication for Node Scheduling. In *Proceedings of the 13th International Conference on Mobile Ad Hoc and Sensor Systems (MASS)* (Brasilia, Brazil). IEEE, 56–64.
- [66] Anna-Brit Schaper. 2019. *Truth be Told: Benchmarking BLE and IEEE 802.15.4*. Master’s thesis. ETH Zurich, Zurich, Switzerland.
- [67] Markus Schuß, Carlo Alberto Boano, and Kay Römer. 2018. 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) (Porto, Portugal). IEEE, 30–35.
- [68] Markus Schuß, Carlo Alberto Boano, and Kay Römer. 2020. Making D-Cube an Open Low-Power Wireless Networking Benchmark. In *Proceedings of the 17th International Conference on Embedded Wireless Systems and Networks (EWSN), poster session* (Lyon, France). Junction Publishing, 164–165.
 - [69] Markus Schuß, Carlo Alberto Boano, Manuel Weber, and Kay Römer. 2017. 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)* (Uppsala, Sweden). Junction Publishing, 54–65.
 - [70] Markus Schuß, Carlo Alberto Boano, Manuel Weber, Matthias Schulz, Matthias Hollick, and Kay Römer. 2019. 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)* (Beijing, China). Junction Publishing, 83–94.
 - [71] Silicon Labs. 2016. EFR32MG12 Gecko Multi-Protocol Wireless SoC Family Data Sheet, Rev. 1.6. [Online] <https://www.silabs.com/documents/public/data-sheets/efr32mg12-datasheet.pdf> – Last accessed: 2021-06-30.
 - [72] Philipp Sommer and Yvonne-Anne Pignolet. 2016. dependable Network Flooding Using Glossy with Channel-Hopping. In *Proceedings of the 17th International Conference on Embedded Wireless Systems and Networks (EWSN), competition session* (Graz, Austria). Junction Publishing, 303.
 - [73] Michael Spörk, Carlo Alberto Boano, and Kay Römer. 2019. 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)* (Catania, Italy). ACM, 7–12.
 - [74] Michael Spörk, Carlo Alberto Boano, and Kay Römer. 2020. Improving the Timeliness of Bluetooth Low Energy in Dynamic RF Environments. *ACM Transactions on Internet of Things* 1, 2 (April 2020).
 - [75] Michael Spörk, Markus Schuß, Carlo Alberto Boano, and Kay Römer. 2021. Ensuring End-to-End Dependability Requirements in Cloud-based Bluetooth Low Energy Applications. In *Proceedings of the 18th International Conference on Embedded Wireless Systems and Networks (EWSN)* (Virtual event). Junction Publishing, 55–66.
 - [76] Aleksandar Stanoev, Adnan Aijaz, Anthony Portelli, and Michael Baddeley. 2020. Closed-Loop Control over Wireless – Remotely Balancing an Inverted Pendulum on Wheels. In *Proceedings of the 17th International Conference on Embedded Wireless Systems and Networks (EWSN), demo session* (Lyon, France). Junction Publishing, 195–197.
 - [77] STMicroelectronics. 2020. STM32WB5MMG Datasheet: Bluetooth Low Energy 5.0 and 802.15.4 Module, DS13252 - Rev 1. [Online] <https://www.st.com/en/microcontrollers-microprocessors/stm32wbxm-modules.html> – Last accessed: 2021-06-30.
 - [78] Felix Sutton, Bernhard Buchli, Jan Beutel, and Lothar Thiele. 2015. Zippy: On-Demand Network Flooding. In *Proceedings of the 13th International Conference on Embedded Networked Sensor Systems (SenSys)* (Seoul, South Korea). ACM, 45–58.
 - [79] Makoto Suzuki, Chun-Hao Liao, Sotaro Ohara, Kyoichi Jinno, and Hiroyuki Morikawa. 2017. Wireless-Transparent Sensing. In *Proceedings of the 14th International Conference on Embedded Wireless Systems and Networks (EWSN)* (Uppsala, Sweden). Junction Publishing, 66–67.
 - [80] Makoto Suzuki, Yasuta Yamashita, and Hiroyuki Morikawa. 2013. Low-Power, End-to-End Reliable Collection Using Glossy for Wireless Sensor Networks. In *Proceedings of the 7th International Vehicular Technology Conference (VTC)* (Dresden, Germany). IEEE, 1–5.
 - [81] Texas Instruments. 2021. CC2652R SimpleLink Multiprotocol 2.4 GHz Wireless MCU datasheet, Rev. H. [Online] <https://www.ti.com/product/CC2652R> – Last accessed: 2021-06-30.
 - [82] Pei Tian, Xiaoyuan Ma, Carlo Alberto Boano, Fengxu Yang, Ye Liu, Xin Tian, Dan Li, and Jianming Wei. 2021. ChirpBox: An Infrastructure-Less LoRa Testbed. In *Proceedings of the 18th International Conference on Embedded Wireless Systems and Networks (EWSN)* (Virtual event). Junction Publishing, 115–126.
 - [83] Davide Vecchia, Pablo Corbalán, Timofei Istomin, and Gian Pietro Picco. 2019. Playing with Fire: Exploring Concurrent Transmissions in Ultra-wideband Radios. In *Proceedings of the 18th International Conference on Sensing, Communication and Networking (SECON)* (Boston, MA, USA). IEEE, 1–9.
 - [84] Andrew J. Viterbi. 1967. Error Bounds for Convolutional Codes and an Asymptotically Optimum Decoding Algorithm. *IEEE Transactions on Information Theory* 13, 2 (April 1967), 260–269.
 - [85] Georg von Zengen, Alexander Baumstark, Alexander Willecke, Ulf Kulau, and Lars Wolf. 2018. How Different Transceiver Hardware Effects Concurrent Transmissions in WSNs. In *Proceedings of the 14th International Conference on Distributed Computing in Sensor Systems (DCOSS)* (New York, NY, USA). IEEE, 139–146.
 - [86] Yin Wang, Yuan He, Dapeng Cheng, Yunhao Liu, and Xiang-yang Li. 2013. TriggerCast: Enabling Wireless Constructive Collisions. In *Proceedings of the 32th International Conference on Computer Communications (INFOCOM)* (Turin, Italy). IEEE, 480–484.
 - [87] Yin Wang, Yunhao Liu, Yuan He, Xiang-Yang Li, and Dapeng Cheng. 2015. Disco: Improving Packet Delivery via Deliberate Synchronized Constructive Interference. *IEEE Transactions on Parallel and Distributed Systems* 26, 3 (March 2015), 713–723.
 - [88] Kamin Whitehouse, Alec Woo, Fred Jiang, Joseph Polastre, and David Culler. 2005. Exploiting The Capture Effect For Collision Detection And Recovery. In *Proceedings of the 2nd International Workshop on Embedded Networked Sensors (EmNetS)* (Sydney, NSW, Australia). IEEE, 45–52.

- [89] Matthias Wilhelm, Vincent Lenders, and Jens B. Schmitt. 2014. On the Reception of Concurrent Transmissions in Wireless Sensor Networks. *IEEE Transactions on Wireless Communications* 13, 12 (Aug. 2014), 6756–6767.
- [90] Dingwen Yuan and Matthias Hollick. 2013. Let’s Talk Together: Understanding Concurrent Transmission in Wireless Sensor Networks. In *Proceedings of the 38th International Conference on Local Computer Networks (LCN)* (Sydney, NSW, Australia). IEEE, 219–227.
- [91] Dingwen Yuan and Matthias Hollick. 2015. 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)* (Boston, MA, USA). IEEE, 1–9.
- [92] Dingwen Yuan, Michael Riecker, and Matthias Hollick. 2014. Making ‘Glossy’ Networks Sparkle: Exploiting Concurrent Transmissions for Energy Efficient, Reliable, Ultra-low Latency Communication in Wireless Control Networks. In *Proceedings of the 11th European Conference on Wireless Sensor Networks (EWSN)* (Oxford, United Kingdom). Springer, 133–149.
- [93] Hui Zhou, Charles Nicholls, Thomas Kunz, and Howard Schwartz. 2008. *Frequency Accuracy & Stability Dependencies of Crystal Oscillators*. Technical Report Technical Report SCE-08-12. Carleton University, Systems and Computer Engineering.
- [94] Marco Zimmerling, Luca Mottola, and Silvia Santini. 2021. Synchronous Transmissions in Low-Power Wireless: A Survey of Communication Protocols and Network Services. *Comput. Surveys* 53 (Feb. 2021), 1–39. Issue 6.