Playing with Fire: Exploring Concurrent Transmissions in Ultra-wideband Radios

Davide Vecchia, Pablo Corbalán, Timofei Istomin, Gian Pietro Picco
University of Trento, Italy
{davide.vecchia, p.corbalanpelegrin, timofei.istomin, gianpietro.picco}@unitn.it

***Abstract*—Concurrent transmissions can be used to enhance network reliability and scalability and to reduce energy consumption and latency. This paper studies their applicability for communication and ranging in UWB networks, where they are hitherto largely unexplored. To this end, we follow an experiment-driven approach and show that *i)* different pulse repetition frequencies virtually double the number of non-interfering channels, *ii)* concurrent transmissions with different preamble codes are unreliable, unless transmitters are tightly synchronized, and *iii)* under the same RF configuration, UWB radios are very likely to receive one of the packets transmitted concurrently by multiple senders, unlocking opportunities similar to those exploited in low-power narrowband radios. We argue that our findings can inform the design of novel communication and ranging protocols exploiting the unique advantages of concurrent transmissions, potentially inspiring a new wave of research on UWB radios.**

I. Introduction

Concurrent Transmissions, i.e., the sending of packets overlapping in space and time, are a fundamental building block of modern protocols for wireless networks.

Concurrent transmissions *over different frequency channels* have been exploited for decades as a simple means to achieve higher scalability w.r.t. users and traffic rates, and avoid interference among senders. For instance, the Time-Slotted Channel Hopping (TSCH) [1], increasingly popular in Internet of Things (IoT) networking, is one of many protocols exploiting the PHY-level separation of IEEE 802.15.4 channels.

In contrast, concurrent transmissions *on the same channel* are commonly considered harmful, as they generally result in packet loss. Nevertheless, this is not always true. Under some tight synchronization requirements, PHY-level radio properties enable reliable packet reception of one of the concurrently transmitted packets, despite interference from the others. This fact has recently been exploited by several protocols yielding unprecedented reliability and performance [2]–[7].

These two nuances of concurrent transmissions are well-studied in the context of mainstream IEEE 802.15.4 narrowband radios (e.g., the popular CC2420 [8]). Nevertheless, very little is reported about them in the context of IEEE 802.15.4 ultra-wideband (UWB) radios—the focus of this paper.

Why UWB? This technology has returned to the forefront of research and market interest after a decade of oblivion, thanks to the recent availability of a new generation of UWB chips significantly smaller, cheaper, and less energy-hungry than their predecessors, spearheaded by the popular DecaWave DW1000 [9]. Moreover, UWB radios can also estimate distance with decimeter-level accuracy. The possibility of using a single radio chip for communication and ranging, and the fact that IEEE 802.15.4 includes an UWB PHY layer, is attracting significant interest in several IoT scenarios.

Concurrent Transmissions in UWB? A large body of work [10], [11] showed that, in theory, the physical encoding of UWB packets allows for non-interfering concurrent transmissions on different links on the same channel, regardless of the power of the signals involved or their synchronization. This is potentially disruptive, as it would mean that properly configured communication links could operate on their own dedicated "virtual" channels, well beyond the limited number of frequency channels—a formidable asset for designing scalable, reliable, and efficient network protocols.

In practice, however, the situation is less clear. The IEEE 802.15.4 standard [1] defines a *complex channel* as composed by a frequency channel and a preamble code, hinting at the fact that different combinations of these two configuration parameters should yield non-interfering communication links. Instead, the documentation of the IEEE 802.15.4-compliant DW1000 states that non-interfering communication links must be obtained by configuring a different pulse repetition frequency (PRF) as with only different preamble codes *"there is still a small amount of cross correlation [. . .] This may mean that it is not possible to achieve the separation envisioned by the standards authors"* (p. 218, [9]). These two guidelines are partially at odds, failing to provide a clear indication to the protocol designer. Further, neither document reports quantitative information about the effect of a given configuration.

More generally, while the real-world characteristics of concurrent transmissions in narrowband radios are well-understood for both the different- and same-channel configurations, this is unfortunately not the case for UWB.

Goals, Methodology, and Contribution. The goal of this paper is precisely to fill this gap by *ascertaining the extent to which concurrent transmissions can be exploited in UWB.*

We focus on communication links in the same collision domain and on the same frequency channel, driven by research questions (§III) aimed at dissecting the conditions under which concurrent transmissions are possible. We specifically investigate the effect of three configurations: *i)* different PRF *ii)* same PRF but different preamble codes *iii)* same PRF and preamble codes. The first two configurations are aimed at obtaining separate, non-interfering complex channels, following the recommendations of DecaWave and IEEE 802.15.4, respectively. Instead, the third one yields interfering links,

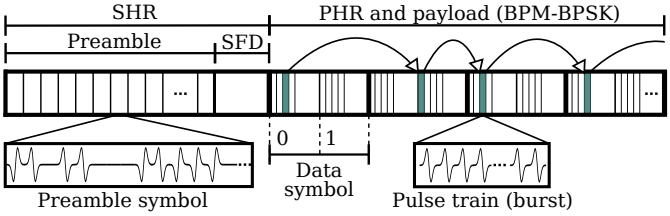

Figure 1: UWB frame: the SHR is encoded in single pulses, while the data part exploits BPM-BPSK modulation. Preamble codes determine the preamble symbol sequence and the time-hopping code (arrows) for data transmission.

enabling us to investigate the presence of PHY-level effects akin to those in narrowband radios [3], [4].

We achieve our goal *experimentally* with a real-world dedicated setup (§IV) where we control precisely the synchronization among nodes and therefore the degree of concurrency, i.e., temporal overlapping among transmitted packets.

Using this setup, we derive empirical observations (§V) on the reliability and performance of *both communication and ranging* for each of the configurations above. To the best of our knowledge, ours is the first study that *quantitatively* investigates both dimensions in the context of concurrent transmissions, and across different notions thereof.

Based on these empirical observations, we distill higher-level findings to *inform the design of networking and ranging protocols* and exemplify opportunities for their application in practice (§VI), hopefully inspiring a new generation of UWB networking and ranging protocol stacks.

The paper ends with a concise survey of related work (§VII) followed by brief concluding remarks (§VIII).

II. Ultra-wideband Communication and Ranging

We provide the necessary background on impulse radio UWB (IR-UWB) and the DecaWave DW1000 transceiver.

Impulse Radio. IR-UWB spreads the signal energy across a very large bandwidth ($\geq$ 500 MHz) by transmitting data through a time-hopping sequence of ns-level pulses [12]. This reduces the power spectral density, the interference produced to other wireless technologies, and the impact of multipath components (MPC). The large bandwidth provides high time resolution, enabling UWB receivers to precisely estimate the time of arrival of a signal, and therefore distance. Time-hopping codes [10] can be used to provide multiple access to the medium. These features make IR-UWB ideal for ranging and localization and also for low-power communication.

UWB PHY Layer. The IEEE 802.15.4-2011 standard [1] specifies an UWB PHY layer based on impulse radio. An UWB frame (Figure 1) is composed of *i)* a synchronization header (SHR) and *ii)* a data portion. The SHR is encoded in single pulses and includes a preamble for synchronization and the start frame delimiter (SFD). The data portion, instead, exploits a combination of burst position modulation (BPM) and binary phase-shift keying (BPSK), and includes a physical header (PHR) and the data payload. The duration of the preamble is configurable and depends on the number of repetitions of a predefined symbol. A preamble symbol (Figure 1) consists of a sequence of elements drawn from a ternary alphabet $\{+1,\ 0, -1\}$, i.e., positive, absent, and negative pulse. This sequence is determined by the preamble code. The standard defines preamble codes of 31 and 127 elements, which are then interleaved with zeros according to a spreading factor. This yields a (mean) *pulse repetition frequency* (PRF) of 16 MHz or 64 MHz, respectively; these values, hereafter $PRF16$ and $PRF64$ for readability, are also configurable. Preamble codes also define the pseudo-random sequence used for time-hopping in the transmission of the data part. Preamble codes were thus envisaged as a mechanism to provide multiple non-interfering access to the wireless medium. However, according to DecaWave [13], frames that overlap in different complex channels ⟨frequency, code⟩ may still interfere with each other unless their codes have different $PRFs$.

Channel Impulse Response (CIR). The perfect periodic autocorrelation of the preamble code sequence enables coherent receivers to determine the CIR [9], which provides information about the multipath propagation characteristics of the wireless channel between a transmitter and a receiver. The CIR allows UWB radios to distinguish the signal's leading edge from MPC and accurately estimate the time of arrival of the signal. In §V, we exploit CIR information to analyze the interference created by concurrent transmissions under different RF configurations.

Two-way Ranging (TWR). The IEEE 802.15.4 standard also specifies two TWR schemes to estimate the distance between two nodes, an initiator and a responder. In the simplest one, single-sided TWR (SS-TWR), the initiator sends a POLL message to the responder, storing the TX timestamp t_1. After an assigned delay δ_{TX}, the responder replies back with a RESPONSE, embedding in the payload the RX timestamp of the POLL, t_2, and the predicted TX timestamp of the RESPONSE, t_3. The initiator then measures t_4, the RESPONSE RX timestamp, computes the time of flight as $\tau = \frac{(t_4-t_1)-(t_3-t_2)}{2}$, and estimates the distance $d = \tau/c$ between the nodes, where c is the speed of light in air. SS-TWR suffers from clock and frequency drift [14]. Symmetric double-sided TWR (DS-TWR), also part of the standard, mitigates their impact but requires more message exchanges. All the timestamps required for ranging are measured in a packet at the *ranging marker* (RMARKER), which marks the first pulse of the PHR after the SFD. In this paper, we focus on SS-TWR and analyze how different RF settings can be exploited to perform multiple ranging exchanges concurrently.

RX Errors. The PHR and data payload employ several mechanisms to detect and correct errors. The PHR includes a 6-bit parity check SECDED (single error correct, double error detect) field. The data payload, instead, employs a Reed-Solomon (RS) encoder that appends 48 parity bits every 330b of data. Hence, uncorrectable bit errors in the PHR or the data payload trigger SECDED and RS errors, respectively. Moreover, UWB radios can also trigger an SFD timeout when a preamble is detected and the SFD is not received within the expected SHR duration. In §V, we analyze these errors to understand the reasons behind packet loss.

DecaWave DW1000. The DW1000 is a standard-compliant UWB transceiver. The DW1000 supports both $PRF16$ and $PRF64$, frequency channels $\{1, 2, 3, 4, 5, 7\}$, and three data rates 110 kbps, 850 kbps, and 6.8 Mbps. Channels $\{4, 7\}$ have a larger 900 MHz bandwidth while the others are limited to 499.2 MHz. The DW1000 measures the CIR with a sampling period of $T_s = 1.016$ ns upon preamble reception, storing it in a large 4096B buffer available to the firmware developer.

III. Research Questions

Concurrent transmissions are a complex, multi-faceted topic. In this section, we concisely outline the intertwined research questions motivating this paper and answered in §V.

Q1: What are the key configuration settings yielding non-interfering complex channels? As already mentioned (§I), the PRF and preamble code values play a key role, but neither the DecaWave documentation nor the IEEE 802.15.4 standard provide an exhaustive answer about how to reliably define complex channels. Therefore we explore both configurations hinted in these documents: *i)* different $PRFs$ and preamble codes, and *ii)* same PRF and different preamble codes.

Answering this question is key, as UWB frequency channels in commercial chips are fewer than supported. For instance, the DW1000 offers only 6 channels, although IEEE 802.15.4 defines 16 channels for both its narrowband and UWB PHY layers. Therefore, using complex channels, instead of only frequency channels, significantly increases the degrees of freedom in scheduling non-interfering concurrent transmissions.

Q2: Can concurrent transmissions be reliably exploited even on the same *complex channel?* To answer this question we experiment also with links configured with the same PRF and preamble codes. This allows us to ascertain whether the recent results [2]–[7] exploiting concurrent transmissions on the same channel in IEEE 802.15.4 narrowband can be transferred or adapted for UWB, and under what conditions.

Q3: How concurrent the concurrent transmissions can be? Or, in other words, what is the tolerable amount of overlapping among transmissions? Answering this question yields precious information to the protocol designer, as it determines the amount of synchronization (or de-synchronization) required to prevent performance degradation.

Q4: Is the outcome affected by the relative power of the concurrent signals? Network nodes are typically configured with the same TX power; nevertheless, the different relative node positions, combined with the well-known path loss attenuation, may induce significant differences in the power of signals concurrently transmitted. This difference in power plays a key role in determining the capture effect in IEEE 802.15.4 narrowband radios [3], [4]; it is therefore worth investigating if similar constraints exist in UWB in the context of Q2. Furthermore, it is also worth investigating the answer of this question in the context of Q1, to ascertain if the relative difference in signal power plays a role in determining non-interference across complex channels.

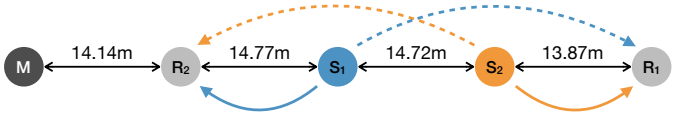

Figure 2: Network topology of our experiments. M, S_i, and R_i are the master, sender, and receiver (responder), respectively. All nodes are in communication range. The arcs represent the links under study: *weak* (dashed line) and *strong* (solid).

Q5: Is the outcome affected by the number of concurrent sources or the transmitted packet? These factors are known to affect concurrent transmissions on the same channel in narrowband IEEE 802.15.4. Therefore, this question relates to Q2 in ascertaining similarities and differences w.r.t. UWB.

We exploit our experimental setup (§IV) to answer these questions via empirical observations (§V) focused on *both communication and ranging* as characterized by the following **metrics**: *i)* packet reception rate (PRR), i.e., the ratio of packets successfully received over those sent *ii)* ranging error (§II), and *iii)* ranging reliability, i.e., the ratio of successful ranging exchanges over those performed.

IV. Experimental Setup

Hardware and Testbed. We run experiments in a testbed deployed in the ceiling above the corridors of an office building at our premises. We employ the DecaWave EVB1000 platform [15], featuring an STM32F105 MCU and the DW1000 UWB transceiver with a PCB antenna. Each EVB1000 is connected to a JTAG programmer and a Raspberry Pi. This setup allows us to easily schedule and run numerous experiments without the effort required to manually deploy the nodes.

Network Topology. Unless otherwise noted, we use a 5-node subset of the testbed (Figure 2) in the same collision domain, with different node roles: *i)* M, synchronization master *ii)* S_i, sender (or SS-TWR initiator) *iii)* R_i, receiver (or responder).

Depending on the senders and receivers chosen, and their relative distance, we can explore different relative signal strengths between transmissions, therefore catering for question Q4. Hereafter, we refer to *strong* links as those where the sender transmit towards its nearest receiver ($S_1 \rightarrow R_2$, $S_2 \rightarrow R_1$) and, dually, *weak* links as those where it transmits to the farthest ($S_1 \rightarrow R_1$, $S_2 \rightarrow R_2$).

Time (De)synchronization. At the start of each message round, the master M broadcasts a synchronization frame. Senders measure the RX timestamp and schedule their transmission after a given delay. Responders log the received packets and RX errors and, for ranging, transmit their RESPONSE.

To assess the impact of time (de)synchronization, and cater for question Q3, we also apply a $\Delta t \in [-183\,\mu s, 183\,\mu s]$ in steps of 32 ns to the transmission of sender S_2, therefore controlling the time overlapping among packets. We checked that reducing the step to the supported minimum of 8 ns does not yield new observations though slows down the experiments. For each time shift, we transmit 25 messages, reporting results from $>$ 296k packets per link and configuration tested. For ranging, we change the time shift range

to $\Delta t \in [-511\,\mu s, 511\,\mu s]$ in steps of 250 ns to account for the overlap of the end of RESPONSE from one link with the beginning of POLL from the other. For each time shift, we perform 10 SS-TWR exchanges, resulting in $\geq$ 39k ranging exchanges per link.

UWB Settings. We consider UWB channels 2 and 4 with center frequency $f_c = 3993.6$ GHz and bandwidth of 499.2 MHz and 900 MHz, respectively. We present results for $PRF16$ and $PRF64$ and with preamble codes $\{3, 4, 9, 10\}$ (channel 2) and $\{7, 8, 17, 18\}$ (channel 4). We set the DW1000 to employ the 6.8 Mbps data rate with a preamble length of 128 symbols.

Implementation. We developed our firmware atop Contiki OS for the EVB1000 platform. At bootstrap, nodes remain idle for 60 s to allow the clock drift to stabilize. To ensure the expected overlapping between frames, we re-synchronize all nodes at the beginning of each round, compensating for the different time propagation between the master M and the other nodes. We leave a 1.5 ms delay between the reception of the synchronization frame and the first transmission round to account for the time required to switch RF configuration and transmit a preamble. Our SS-TWR implementation sets the RESPONSE delay to $\delta_{TX} = 320\,\mu s$ to minimize the impact of clock drift on ranging estimation. In our testbed, we measured a typical clock drift ≤ 3 ppm; this yields a potential desynchronization up to 4.5 ns, negligible as it is < 8 ns, the DW1000 TX scheduling precision [9].

V. Empirical Observations

We present our empirical observations, aimed at answering the research questions in §III using the setup in §IV. We first establish a baseline for communication and ranging by analyzing their performance with isolated transmissions (§V-A). We then structure the core of this section around the configurations we explore concurrent transmissions with: *i)* different $PRFs$ (§V-B) *ii)* different preamble codes within the same PRF (§V-C) *iii)* exact same RF settings (§V-D). Finally, we investigate their combination (§V-E).

A. Baseline: Isolated Transmissions

We first establish the baseline performance of each link *in isolation*, i.e., without concurrent transmissions. The results of communication experiments are averaged over 10000 packets sent by each sender. Across all experiments, we obtain $PRR \geq 99.92\%$ and $PRR \geq 99.99\%$ for $PRF16$ and $PRF64$, respectively. When the sender is closest to the expected receiver (strong links) both $PRFs$ achieve 100%. For ranging, we calibrate the antenna delay for all configurations, obtaining zero-mean error with standard deviation $\sigma \leq 4.5$ cm.

B. Concurrent Transmissions with Different PRFs

Communication Reliability. Figure 3 shows the PRR on channel 4 for each link and PRF combination vs. the applied time shift. Overall, with $PRF64$ we obtain $PRR \geq 96\%$ irrespective of the time shift, while $PRF16$ achieves a slightly lower $PRR \geq 88\%$. When the sender is farther from the intended receiver (Figure 3a), only 11 packets were lost out of

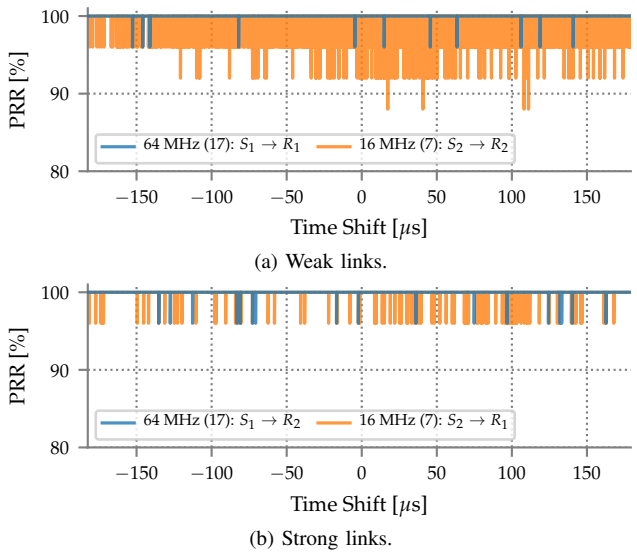

(a) Weak links.

(b) Strong links.

Figure 3: PRR with different $PRFs$ (channel 4). Concurrent transmissions exploiting different PRF are very likely to be received correctly, especially with $PRF64$.

296800 with $PRF64$, yielding an average $PRR = 99.996\%$ despite the concurrent transmissions with $PRF16$. The latter lost 1722 packets, yielding $PRR = 99.42\%$. When the sender is closer to the receiver (Figure 3b), only 17 and 103 packets were lost for $PRF64$ and $PRF16$, respectively, yielding $PRR \geq 99.97\%$. We obtain similar results with other preamble code combinations on channel 2 and 4. The higher reliability of $PRF64$ is the result of the higher amount of pulses per preamble symbol [16]. This comes, however, at a cost in terms of energy. The few packet losses obtained are mostly the result of SECDED and RS errors, i.e., non-correctable bit errors in the PHY header or in the payload. Overall, we observe that concurrent transmissions through different $PRFs$ are reliable regardless of the time (de)synchronization and the physical arrangement of the network.

Ranging Reliability and Error. Figure 4 shows the ranging reliability (top) and the ranging error (bottom) over the applied time shifts for the two weak links ($S_1 \to R_1$ and $S_2 \to R_2$) on channel 4. With $PRF64$, we obtain an average ranging error $\mu = 1$ cm with standard deviation $\sigma = 5$ cm. $PRF16$ yields $\mu = 0.2$ cm with $\sigma = 3$ cm. These results are in accordance with the baseline in isolation (§V-A), indicating that performing ranging concurrently with two different $PRFs$ has no impact on accuracy. With $PRF16$, however, we observe a minor overestimation of the ranging distance for time shifts $\Delta t \in [-50\,\mu s, 50\,\mu s]$. On channel 2, $PRF16$ also presents some extreme, although rare, outliers in the same region. The error was > 1 m for only 0.04% of the ranging samples. For both channels and $PRFs$, we obtain a ranging reliability $\geq 98.27\%$. $PRF16$ is slightly less reliable than $PRF64$. The minimum reliability for a given time shift was 60% with $PRF64$ and channel 2. The minor loss in reliability w.r.t. the PRR in the communication experiment is expected as each SS-TWR exchange requires *two* packets. Figure 5 shows the

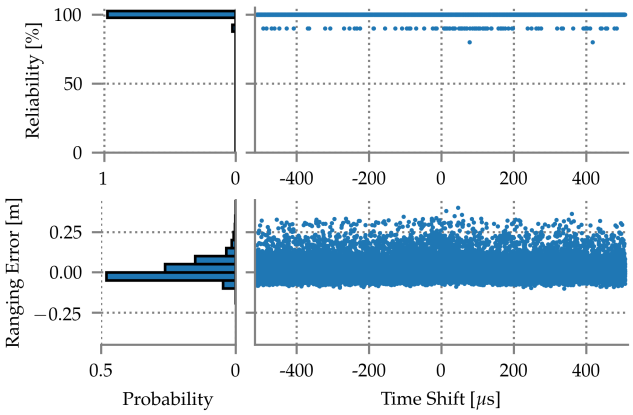

(a) Channel 4, $PRF64$, code 17: $S_1 \rightarrow R_1$.

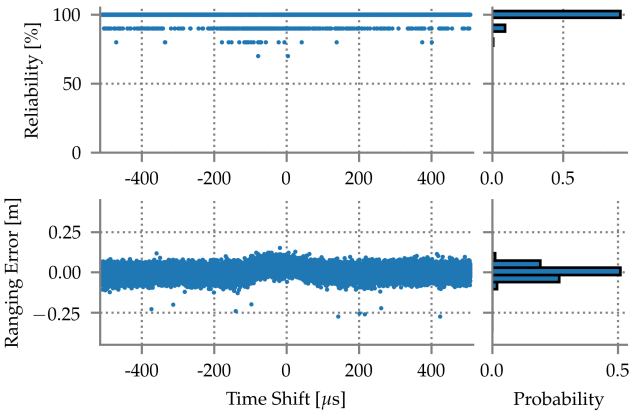

(b) Channel 4, $PRF16$, code 7: $S_2 \rightarrow R_2$.

Figure 4: Concurrent ranging with different $PRFs$. Despite interference, both $PRFs$ perform accurate ranging reliably.

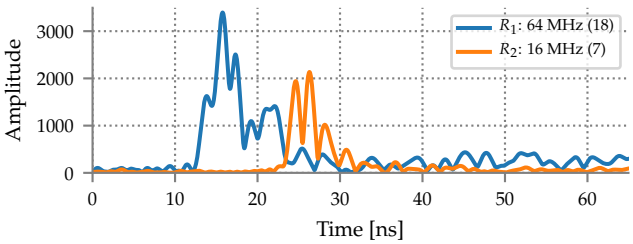

Figure 5: CIR with concurrent transmissions using different $PRFs$ and $\Delta t = -20.513\,\mu s$.

CIRs measured by R_1 and R_2 with both senders transmitting concurrently and $\Delta t = -20\,\mu s$. Both CIRs exhibit a clear line-of-sight path followed by some strong MPC, without impact from concurrent transmissions. As a result, both receivers can distinguish the first path and measure distance accurately.

C. Concurrent Transmissions with Different Preamble Codes

Communication Reliability. Figure 6 and 7 show the PRR of each *weak* and *strong* pair, respectively, for each applied time shift. In contrast to the case with different $PRFs$, concurrent transmissions with different preamble codes introduce significant packet loss, decreasing reliability across the time shifts applied and RF settings studied to $42\% \leq PRR \leq 53\%$. We observe, however, that the *early* packet is likely to be successfully received at the intended destination, especially if the packet is sent $\geq 100\,\mu s$ earlier that the interfering packet. In this case, the end of the preamble or the data portion of the early packet overlaps with the beginning of the preamble of the late packet, bearing reduced impact in the successful reception of the first. This observation could be exploited, e.g., to give priorities to different packets, allowing high priority transmissions to start sufficiently early. The low cross-correlation between preamble codes allows both receivers to synchronize with the early preamble, decreasing the probability of reception for the late packet.

When the sender is close to the receiver and the signal is stronger (Figure 7), transmitting synchronously with the interfering signal or slightly earlier provides high reliability, which underlines the importance of the relative signal strength among concurrent transmissions (Q4). If the interfering signal is weaker and frames are precisely synchronized, we obtain high PRR for each link. With $PRF64$ and $\Delta t < 10\,\mu s$, the overall PRR was 99.81%. We noticed a clear asymmetry (Figure 7b) between the two links; packets transmitted from sender S_2 are more likely to be received than those from S_1. This is the result of the slightly different distance among nodes (Figure 2); we verified it by temporarily moving nodes to the same distance, obtaining more symmetric performance.

CIR Analysis. We resort to the measured CIRs to understand the reasons behind the performance degradation w.r.t. the case with different $PRFs$. Figure 8 shows the CIR estimated by each receiver for various time shifts. When S_2 transmits sufficiently early (Figure 8a), the intended receiver (R_2) receives numerous preamble symbols without any interference, accumulating enough energy for the line-of-sight peak to emerge from other minor MPC and noise. The same occurs, reversed, with link $S_1 \rightarrow R_1$ (Figure 8c). In these cases, the early preamble can be easily detected and the packet is likely to be received correctly. The late transmission, however, suffers strongly from interference of the other, yielding minor peaks throughout the entire CIR span that hinder precise synchronization and first-path estimation. This effect is exacerbated in Figure 8b, where transmissions are more synchronized and it is more difficult to discern the right peak.

Understanding Packet Loss. Table I reports results from an isolated link we misconfigured to use different preamble codes for sender and receiver. As expected, no packet was received correctly, but with counter-intuitive sources of error (§II).

Given the different codes and CIR signals (Figure 8) we expected mostly SFD timeouts or no preamble detection. Instead, 99% of the errors (with either PRF) are due to SECDED or RS parity checks, i.e., non-correctable bit errors in the PHR and data payload, respec-

Table I: RX errors on a misconfigured link.

Error	PRF16	PRF64
SFD t/out	0.28%	0.03%
SECDED	81.55%	86.45%
RS	18%	13.29%
CRC	0.16%	0.22%

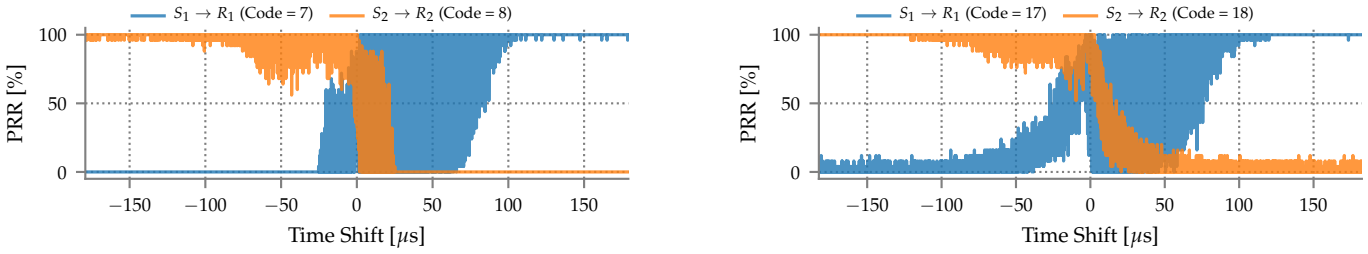

(a) $PRF16$. (b) $PRF64$.

Figure 6: PRR on the *weak* links (channel 4) and different preamble codes for $PRF16$ (left) and $PRF64$ (right). Concurrent transmissions with different preamble codes introduce significant packet loss, especially for the late transmission.

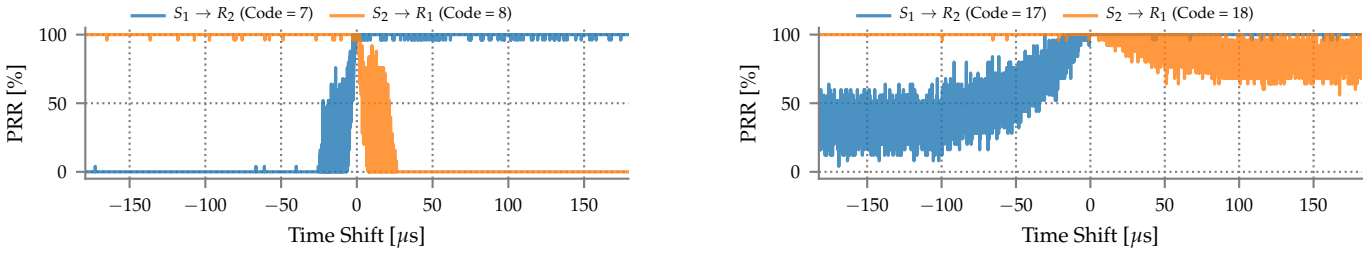

(a) $PRF16$. (b) $PRF64$.

Figure 7: PRR on the *strong* links (channel 4) and different preamble codes for $PRF16$ (left) and $PRF64$ (right).

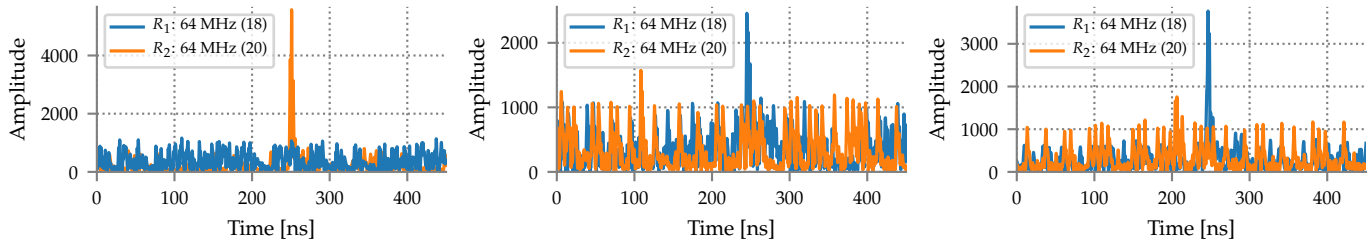

(a) $\Delta t = -41\,\mu s$ (b) $\Delta t = 10\,\mu s$ (c) $\Delta t = 41\,\mu s$

Figure 8: CIR for various time shifts Δt with concurrent transmissions using different preamble codes.

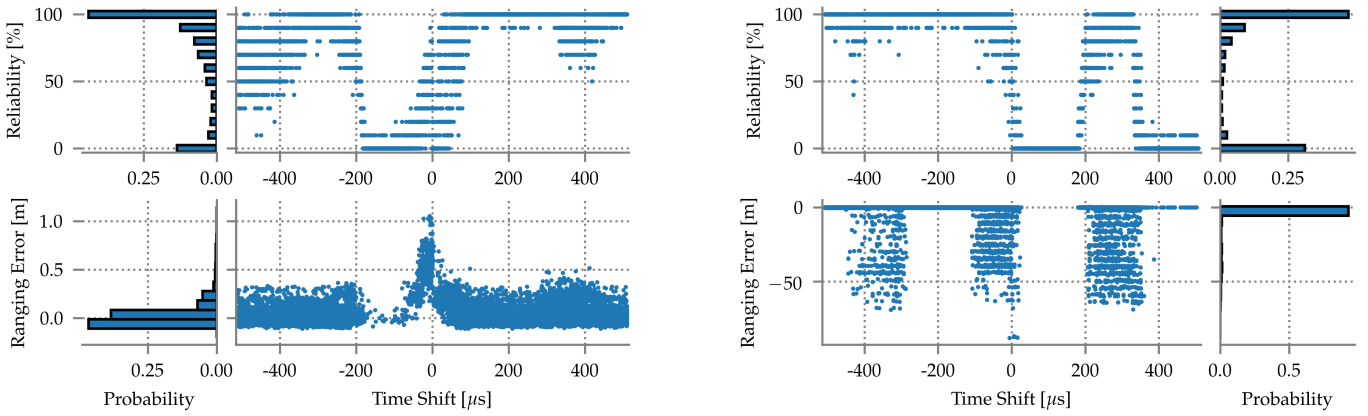

(a) Channel 4, $PRF64$, code 17: $S_1 \to R_1$. (b) Channel 4, $PRF16$, code 8: $S_2 \to R_2$.

Figure 9: Concurrent ranging with different preamble codes. Significant outliers appear especially with $PRF16$. Many ranging rounds are lost due to interference and RX errors.

tively. This suggests that, despite the low cross-correlation of preamble codes, the receiver synchronizes to the preamble sent with the different code and is even able to detect the SFD. If preamble codes were fully orthogonal, a receiver would be able to distinguish the different preamble codes, rejecting or ignoring the mistaken preamble.

Ranging Reliability and Error. From the results on communication, it descends that many ranging rounds cannot be

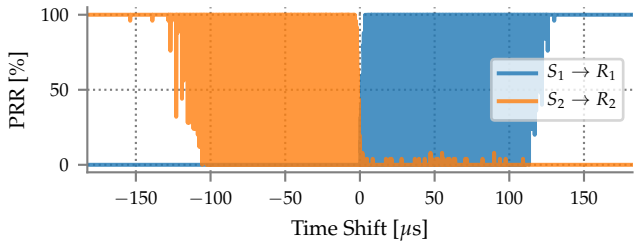

(a) *Weak* links: the intended signal is weaker than the interfering one.

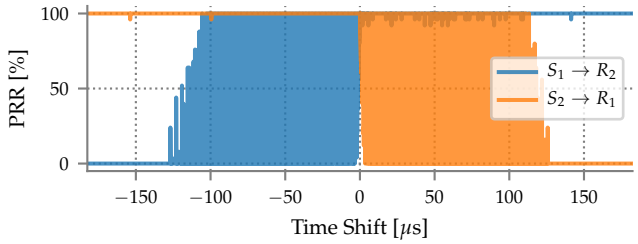

(b) *Strong* links: the intended signal is stronger than the interfering one.

Figure 10: PRR with the same RF configuration (channel 4, $PRF64$, code 17).

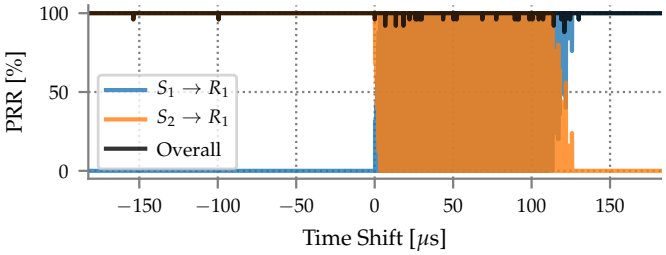

Figure 11: PRR in a single-receiver scenario (channel 4, $PRF64$, code 17).

completed because nodes detect the first preamble transmitted, even if it is the one used by the other link. Figure 9 illustrates the reliability of ranging rounds and the measurement error, depending on the time shift. The success rate follows similar patterns for all configurations. Note that initiators transmit the POLL, switch off their radio, and wake up just in time for the expected RESPONSE. For $S_1 \to R_1$ (Figure 9a), when S_2 initiates the ranging exchange earlier than S_1, the responder R_1 misses the POLL of S_1 because it is receiving the preamble of the POLL from S_2. This mismatch causes the failure of rounds in the range $[-200\,\mu s, 0\,\mu s]$, where the responder cannot recover from the RX error fast enough to receive its intended packet. This behavior is mirrored for S_2 (Figure 9b) in $[0\,\mu s, 200\,\mu s]$. The other relevant drops in reliability are caused by similar interactions w.r.t. RESPONSE.

As for ranging error, $PRF64$ generally achieves accurate distance estimates. Instead, $PRF16$ yields a meter-level error standard deviation, due to the magnitude of outliers whose position is nonetheless well-delimited (Figure 9b). The time shifts associated to outliers are those for which the SFDs of the link at stake overlap with the preamble of another frame, including both POLL-POLL and POLL-RESPONSE conflicts.

D. Concurrent Transmissions with the Same RF Configuration

Communication Reliability. In this case, each receiver might get either the frame sent by the intended sender, the competing one addressed to the other receiver, or none if collision occurs.

Figure 10 shows results for both weak and strong links, with the overall average PRR remaining at 45% and 54.6% respectively. The charts are symmetric w.r.t. the zero-shift axis, reflecting the fact that the links are equivalent. When S_2 transmits earlier (left side of the charts), its packet is likely to be received by the intended target. If its signal is stronger than the interfering one (Figure 10b), its packet is received with nearly 100% probability. Instead, when the delayed interfering signal is stronger (Figure 10a), the PRR for S_2 covers the whole 0–100% range, achieving $\sim$ 90% on average.

Same Receiver. When studying the areas of the charts where the PRR fluctuates, we noticed that they complement each other, i.e., when a receiver misses a frame from its intended sender, it likely receives the interfering frame instead. To see it clearly, we visualize the data differently (Figure 11) by focusing on a single receiver R_1 and plotting, for every time shift, the amount of packets R_1 receives from either S_1 or S_2, and the total. First, we note that the overall PRR remains $\sim$ 100% throughout the tested range, witnessing a very low rate of collisions in a situation with two transmitters competing on the medium to reach the same receiver. Further, we confirm that if the early signal is stronger ($S_2 \to R_1$) this is the one received (left side). Instead, when it is weaker ($S_1 \to R_1$) the radio often "switches" to the stronger one (right side) when it comes during the preamble of the weaker one.

As visible in Figure 12, this switching occurs roughly every 1 µs and lasts for ∼136 ns. Its periodicity matches the duration of a single preamble symbol, suggesting that the radio is able to ignore the stronger frame and keep receiving the weaker one if the symbols of the two preambles are displaced enough. Otherwise, if they coincide, the radio switches to the stronger frame. Interestingly, there are no fluctuations when the absolute shift exceeds 140 µs, i.e., when the later (even if stronger) preamble arrives after the SFD of the first frame was received.

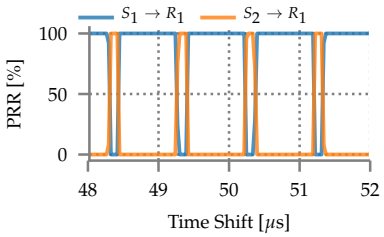

Figure 12: Zoom-in of Fig. 11.

The Role of SFD Timeouts. We ran these experiments by configuring the radio SFD timeout to be larger than the duration of two preambles. A lower value causes significant packet loss as the radio often *i)* synchronizes with the weaker preamble and starts accumulating preamble symbols *ii)* switches to the stronger delayed preamble *iii)* misses the weak SFD because it is "overridden" by the stronger preamble, and *iv)* eventually times out before having a chance to receive the stronger SFD. Therefore, the SFD timeout is crucial for concurrent transmissions when all nodes share the same RF configuration. However, we verified that it bears no influence in the previous cases with different PRF and/or preamble code.

Multiple Transmitters. As noted in §I, many modern protocols for low-power wireless networks build on tightly syn-

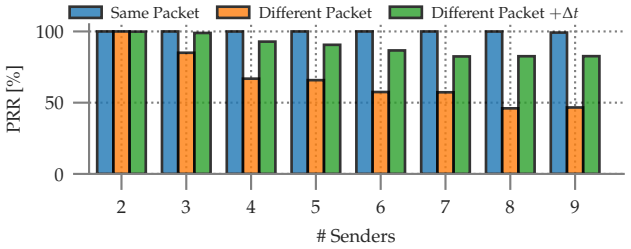

Figure 13: PRR vs. number of concurrent senders (ch.4, $PRF64$, code 17).

chronized network floods. In these protocols, multiple nodes may transmit the same packet (or even different ones) concurrently with the receivers being able to reliably decode one of them. Therefore, we study the performance of UWB receivers when multiple transmitters broadcast concurrently, and extend our experimental setup with more nodes. One of them was configured as the unique receiver, while up to 9 others served as concurrent senders; we ensured they remained tightly synchronized. The PRR is computed based on successful reception of *any* of the concurrent packets.

First, we looked at the worst case where all senders are arranged in a circle, 2 m away from the receiver. In this scenario, almost no communication occurred when senders used *different* packets. This was expected because all arriving signals had similar strengths and timing, and therefore the receiver was unable to discern them. However, when all senders used *identical* packets, this same-distance network yielded highly variable results. The choice of senders (among the 9 available) and their relative distance affected the results tremendously, with the PRR varying in 0–100% even with only 2 senders. This indicates that slight variations in signal propagation paths, indoor reflections, and manufacturing differences among the nodes (e.g., clock drift, radiated power) may cause destructive interference.

In practical network deployments, however, the distances among nodes are never exactly the same. Thus, we repeated the experiments by placing the nodes along a corridor of our building. We did *not* compensate the signal propagation delay as in the aforementioned protocols all neighbors are potential receivers, with different distances.

Figure 13 shows the PRR as we add more concurrent transmitters in sequence, starting from the two closest to the receiver to the farthest one. When the whole network transmits the same packet, it is received in >99% of the cases. When different packets are transmitted synchronously, instead, PRR decreases as the senders increase, down to $< 50\%$ for 9 senders. Interestingly, the results are better ($PRR > 85\%$) when the transmissions are intentionally scattered by adding a random jitter within 20 µs; even if packets are different, the time gap between them enables the receiver to synchronize with the first one and stick to it till the end of the reception. $PRF16$ shows the same relative trends but with worse absolute values, e.g., $PRR < 35\%$ for 9 senders synchronously transmitting different packets and 81% with the jitter added.

Ranging. Although one of the concurrently transmitted packets is likely to be received, the inherent non-determinism severely undermines ranging as nodes may receive, e.g., the RESPONSE from the wrong responder. This observation is reported in [17], which explores the feasibility of a concurrent ranging scheme where the individual "peaks" from simultaneous responders are recovered from a single CIR at the initiator.

E. Combined Settings

We showed (§V-B) that links with different $PRFs$ barely affect each other. We also showed (§V-D) that one of the packets from multiple senders with the same RF settings can be received with high probability. The question is whether the two properties can be exploited at the same time; this would enable, e.g., to run non-interfering instances of the same flood-based protocol over different complex channels.

Table II: PRR for 2 colocated networks with different $PRFs$.

#senders	PRF16	PRF64
one network active (baseline)		
9	–	99.98
9	98.35	–
both networks active		
2+2	96.08	99.98
3+3	95.41	99.92
5+5	84.65	99.96
9+9	86.30	99.20

To this end, we doubled the 10-node deployment in §V-D by pairing each of its nodes with another one configured to use the same frequency channel (4) but a different PRF. We also synchronized the two networks ensuring that all nodes transmit the same packets simultaneously.

Table II shows the PRR for the two receivers in each network vs. varying number of senders, compared to the baseline obtained with the two networks isolated. Although the reliability of the complex channel with $PRF16$ is clearly affected by the increase in senders, it nonetheless remains >84%, making it a useful design choice. On the other hand, this setup confirms the reliability of $PRF64$, yielding a $PRR > 99\%$ even under the heaviest load.

VI. Discussion

Table III summarizes our findings. In contrast to DecaWave's claims (p. 15, [13]), we found that there is some interference between different $PRFs$. $PRF16$ is more affected, especially when the interferer is close to the receiver. In practice, complex channels with different $PRFs$ are almost independent for both communication and ranging, and can be used to deploy co-located yet separate networks (e.g., to increase the scalability of localization systems), or to enhance parallelism in multi-channel protocols like TSCH.

Different preamble codes with the same PRF, unfortunately, do not provide independent channels due to cross-code interference. The rate of collisions can be significantly reduced by synchronizing well the senders, and communication was highly reliable when the interferer was farther than the sender. Arguably, these properties are likely to find application only in very specific, niche application cases.

Finally, our tests with concurrent transmissions in the same complex channel yielded very positive outcome for protocols relying on synchronous transmissions or contention-based medium access. Indeed, the very successful reception (99%) of simultaneous transmissions of the *same* packet from multiple

Table III: Summary of findings.

	Communication	Ranging
Different PRFs	Concurrency is possible with high reliability for both $PRFs$. $PRF16$ is slightly affected by the interference, showing minor packet loss.	Ranging is reliable and measurements are precise. In channel 2, using $PRF16$ results in some (rare) outliers.
Same PRF, different codes	Preamble codes do not provide independent channels. Only the first frame sent is likely to be received. Reliability increases if frames are synchronized, especially when the interference source is far from the destination.	Ranging exchanges may fail or give imprecise estimates depending on how POLL and RESPONSE overlap. $PRF16$ is susceptible to extreme outliers.
Same PRF and code	One of the concurrent frames is likely to be received regardless of the way they overlap. If payloads are identical, sending them synchronously ensures near-perfect reliability. Instead, different packets must not be precisely synchronized to avoid collisions. SFD timeout should be increased.	Similarly to the case with different preamble codes, ranging exchanges may fail depending on the time shift between initiators. Both $PRF64$ and $PRF16$ are affected by extreme outliers.

nodes in principle enables techniques like those in A-MAC [2], Glossy [4] and others [5] on UWB. Further, the fact that *different* packets rarely collide destructively if shifted by a tiny jitter enables the techniques in Crystal [6] and Chaos [7].

VII. Related work

Early work on UWB investigated, mostly theoretically, methods enabling multiple access to the wireless medium, e.g., different time-hopping codes [10], [11] or orthogonal pulse shapes [18], [19]. The IEEE 802.15.4 standard [1] is based on the former although, as shown in §V, it results in unreliable performance unless codes use different $PRFs$, as observed by DecaWave [13]. Our work goes beyond this observation, and offers quantitative evidence from real-world experiments about the expected performance using different $PRFs$ and codes, informing the design of communication and ranging schemes, to seize the opportunities offered by concurrent transmissions.

Few works investigated these opportunities. SurePoint [20] employs a Glossy-like flooding primitive to schedule ranging exchanges between mobile tags and anchors; however, it is neither detailed nor evaluated as it is not the focus of the paper. Concurrent ranging [17] exploits tightly-synchronized transmissions to measure the distance to several devices on a single TWR exchange by analyzing the CIR signal information, offering experimental results geared towards the specific technique proposed. Neither work analyzes the performance of concurrent transmissions w.r.t. different $PRFs$ and preamble codes, nor evaluates the impact of time (de)synchronization or different signal power—key aspects of general applicability, and addressed by our paper.

Our experimental analysis is inspired by the more established work on concurrent transmissions in low-power narrowband radios with several protocols providing efficient and reliable multi-hop communication. These protocols exploit, in the same RF channel, the PHY-level properties that allow the radio successfully decode data when multiple senders transmit identical or even different packets simultaneously. Our results in §V-D suggest that similar benefits can be seized in UWB, potentially inspiring a new wave of research on UWB communication protocols based on concurrent transmissions.

VIII. Conclusions

Concurrent transmissions are a powerful tool for the designers of communication protocols, and have been successfully exploited in different ways by many works in IEEE 802.15.4 narrowband radios. Unfortunately, the same does not hold for UWB radios, whose peculiar ability to combine high-rate communication and accurate distance estimation is placing them at the forefront of IoT scenarios. Indeed, the guidelines in the IEEE 802.15.4 standard conflict with the recommendations for the most popular UWB chip, the DW1000.

We analyzed the conditions under which concurrent transmissions can be reliably exploited under different radio settings, to a depth and extent hitherto unreported in the literature. The high-level findings we distilled can be immediately exploited by protocol designer, potentially inspiring a new generation of UWB networking and ranging protocols exploiting the advantages of concurrent transmissions.

References

[1] IEEE 802.15.4-2015, Standard for Low-Rate Wireless Networks.

[2] P. Dutta, S. Dawson-Haggerty, Y. Chen, C. J. M. Liang, and A. Terzis. Design and Evaluation of a Versatile and Efficient Receiver-initiated Link Layer for Low-power Wireless. In *Proc. of SenSys*, 2010.

[3] J. Lu and K. Whitehouse. Flash Flooding: Exploiting the Capture Effect for Rapid Flooding in Wireless Sensor Networks. In *Proc. of INFOCOM*, 2009.

[4] F. Ferrari et al. Efficient Network Flooding and Time Synchronization with Glossy. In *Proc. of IPSN*, 2011.

[5] F. Ferrari, M. Zimmerling, L. Mottola, and L. Thiele. Low-power Wireless Bus. In *Proc. of SenSys*, 2012.

[6] T. Istomin et al. Data Prediction + Synchronous Transmissions = Ultra-low Power Wireless Sensor Networks. In *Proc. of SenSys*, 2016.

[7] O. Landsiedel et al. Chaos: Versatile and Efficient All-to-all Data Sharing and In-network Processing at Scale. In *Proc. of SenSys*, 2013.

[8] Texas Instruments. CC2420 Datasheet.

[9] DecaWave Ltd. DW1000 User Manual, version 2.15, 2017.

[10] M. Z. Win and R. A. Scholtz. Ultra-Wide Bandwidth Time-Hopping Spread-Spectrum Impulse Radio for Wireless Multiple-Access Communications. *IEEE Transactions on Communications*, 48(4):679–689, 2000.

[11] F. Ramirez-Mireles and R. A. Scholtz. Multiple-Access Performance Limits with Time Hopping and Pulse Position Modulation. In *Proc. of the IEEE Military Communications Conference (MILCOM)*, 1998.

[12] M. Z. Win and R. A. Scholtz. Impulse radio: How it works. *IEEE Comm. Lett.*, 2(2):36–38, 1998.

[13] DecaWave Ltd. APH010 DW1000 Inter-channel Interference: How transmissions on one DW1000 Channel can affect other channels and how to minimize that effect. Version 1.1., 2017.

[14] DecaWave Ltd. APS011 Application Note: Sources of Error in DW1000 based Two-way Ranging (TWR) Schemes, 2014.

[15] DecaWave Ltd. DecaWave ScenSor EVB1000 Evaluation Board, 2013.

[16] B. Großwindhager, C. Alberto Boano, M. Rath, and K. Römer. Enabling Runtime Adaptation of Physical Layer Settings for Dependable UWB Communications. In *Proc. of WoWMoM*, 2018.

[17] P. Corbalán and G. P. Picco. Concurrent Ranging in Ultra-wideband Radios: Experimental Evidence, Challenges, and Opportunities. In *Proc. of EWSN*, 2018.

[18] B. Parr et al. A Novel Ultra-Wideband Pulse Design Algorithm. *IEEE Communications Letters*, 7(5):219–221, 2003.

[19] M. Ghavami et al. Novel UWB Pulse Shape Modulation System. *Wireless Personal Communications*, 23(1):105–120, 2002.

[20] B. Kempke et al. SurePoint: Exploiting Ultra Wideband Flooding and Diversity to Provide Robust, Scalable, High-Fidelity Indoor Localization. In *Proc. of SenSys*, 2016.