

Concurrent Ranging in Ultra-wideband Radios: Experimental Evidence, Challenges, and Opportunities

Pablo Corbalán
University of Trento, Italy
p.corbalanpelegrin@unitn.it

Gian Pietro Picco
University of Trento, Italy
gianpietro.picco@unitn.it

Abstract

This paper investigates the feasibility of a novel primitive for *concurrent ranging* in ultra-wideband (UWB) radios. Conventional ranging protocols schedule the packet transmissions used to estimate distance to be separate in time. In contrast, our concurrent ranging primitive relies on the *overlapping* of these transmissions; when a ranging request is issued by an initiator node, all nodes in range immediately reply back. These concurrent signals are “fused” in the communication channel, whose channel impulse response (CIR) is made available by the DecaWave DW1000 transceiver we use in this paper. Combined with the fact that UWB transmissions rely on very short (< 2 ns) pulses, the CIR enables the initiator to determine the precise timing of the individual signals, and therefore estimate accurately the distance of the corresponding responders.

Concurrent ranging removes the need for scheduling, and is faster and less energy-hungry than traditional scheduled schemes. We report empirical evidence that our idea is feasible, and use it as a basis to elicit the challenges and limitations of concurrent ranging, the opportunities for significant improvements w.r.t. the current state of the art, and the preliminary techniques that can be used to achieve them.

Categories and Subject Descriptors

C.2.1 [Computer-Communication Networks]: Wireless Communication

General Terms

Design, Experimentation, Measurement, Performance

Keywords

Ultra-wideband, distance estimation, concurrent ranging

1 Introduction

The recent availability of tiny, low-power ultra-wideband (UWB) radios has renewed interest in this technology, whose

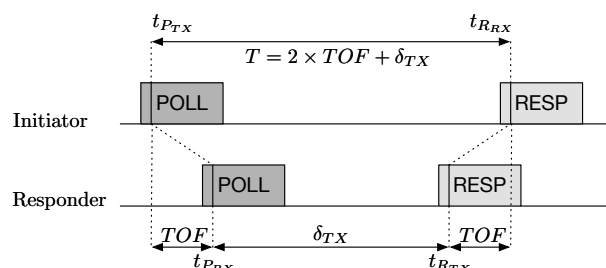


Figure 1. Single-sided two-way ranging (SS-TWR).

peculiarity is to enable accurate distance estimation (*ranging*) along with high-rate communication. UWB is therefore becoming a key player in the Internet of Things (IoT), by providing two key functionalities with a single radio chip.

The Decawave DW1000 transceiver [1] is at the forefront of this new trend, as it provides centimeter-level ranging accuracy with a tiny form factor and a power consumption an order of magnitude lower than its bulky UWB predecessors. On the other hand, its consumption is still an order of magnitude higher than other IoT low-power wireless radios. For instance, a packet transmission requires ~ 300 mW on the DW1000 and is typically < 40 mW on BLE transceivers. Moreover, this higher consumption is amplified when ranging—the key asset of UWB—is exploited.

Ranging: Basics. Conventional ranging protocols are based on pairwise packet exchanges between two nodes: an *initiator* and a *responder*. For instance, Figure 1 illustrates single-sided two-way ranging (SS-TWR), the simplest scheme, also part of the IEEE 802.15.4-2011 standard¹ [2]. In SS-TWR, the initiator requests a ranging measurement via a POLL packet, to which the responder replies with a RESPONSE packet. The RESPONSE packet contains the timestamps marking the receipt of POLL and the sending of RESPONSE on the responder. This information, along with the dual timestamps marking the sending of POLL and the receipt of RESPONSE at the initiator, enable the latter to accurately compute the time of flight

$$t_{TOF} = \frac{t_{R_{RX}} - t_{P_{TX}} - (t_{R_{TX}} - t_{P_{RX}})}{2} \quad (1)$$

and consequently estimate the distance from the responder as $d = t_{TOF} \times c$, where c is the speed of light in air.

¹Hereafter, we follow the terminology used by the DecaWave documentation; the IEEE standard uses *originator* instead of *initiator*.

Motivation. The key point is that two-way ranging, as the name suggests, involves a pairwise exchange between the initiator and every responder. In other words, if the initiator must estimate its distance w.r.t. N nodes, $2 \times N$ packets are required. The situation is even worse with other schemes that improve accuracy by acquiring more timestamps via additional packet transmissions, e.g., $4 \times N$ in the popular double-sided two-way ranging (DS-TWR). Other variations (e.g., as proposed in [3]) ameliorate the situation by *i*) replacing the unicast POLL with a broadcast one, and *ii*) using different times δ_{TX} for each responder, therefore scattering their RESPONSE packets to avoid collisions. This reduces the number of packets to $N + 1$ in SS-TWR. However, it still requires knowledge of the initiator’s neighborhood to determine which responders are available and how to properly schedule their responses—something necessary also in the approaches above. This is a one-time cost for applications in which the set of nodes targeted by ranging is known and fixed throughout the execution; it is a major source of overhead in others where this set is continuously changing, as in applications involving mobile nodes ranging among themselves or moving in a large area where the fixed ranging targets (anchors) change dynamically.

Key idea: Concurrent ranging. We propose a novel approach to ranging in which, instead of *separating* the pairwise exchanges necessary to ranging, these are *overlapping* in time. This *concurrent ranging* primitive requires only $N + 1$ packets *without* any need for scheduling and/or neighbor discovery, and is extremely simple in its mechanics: when a single POLL sent by the initiator is received, each responder sends back its RESPONSE *immediately*, i.e., without any scheduling across nodes to avoid collisions. This primitive significantly reduces latency, which becomes equal to the duration of a single ranging exchange. However, the result is that these concurrent signals from different responders are “fused” in the communication channel, which indeed may lead to a collision at the initiator.

This is precisely where the peculiarities of UWB communications come into play. As we discuss in Section 3, UWB transmissions rely on very short (< 2 ns) pulses, which enable a very precise timestamping of an incoming radio signal. This is what makes UWB intrinsically more amenable to accurate ranging than narrowband, whose reliance on carrier waves that are more “spread” in time induces physical bounds on the precision that can be attained in establishing a time reference for an incoming signal.

These considerations have implications also in the context of our proposed concurrent ranging. In narrowband, the fact that concurrent signals are spread over time makes them also very difficult to tell apart once fused into a single signal. Further, this is possible only if detailed channel state information is available—usually not the case on narrowband low-power radios, e.g., the popular CC2420 and its recent descendants.

In contrast, UWB reliance on short pulses makes concurrent signals less likely to overlap therefore enabling, under certain conditions discussed later, their identification if channel state information is available. Interestingly, the DW1000 *i*) bases its own operation precisely on the processing of the channel impulse response (CIR), and *ii*) makes the CIR avail-

able also to the application layer.

Goals and contributions. In principle, concurrent ranging removes the need for scheduling and neighbor discovery, is faster and less energy-hungry than traditional scheduled schemes, and therefore holds the potential for a breakthrough in the field of UWB ranging. Nevertheless, the actual feasibility of this idea has hitherto not been demonstrated empirically or explored theoretically; to the best of our knowledge we are the first to propose it, as discussed in Section 2.

The first and foremost goal and contribution of this paper is therefore to ascertain the *feasibility* of concurrent ranging, for which we adopt an experimental approach based on the small-scale setup described in Section 6. Specifically, we investigate the inevitable degradation in accuracy w.r.t. isolated ranging, due to the interference among the responder’s signals, and explore the extent to which the latter constrains the minimum and maximum distance among nodes beyond which it is impossible to separate these signals. Our results, discussed in Section 6, offer empirical evidence that it is possible to derive accurate ranging information from UWB signals overlapping in time.

More exhaustive measurements are needed to corroborate these findings, e.g., in larger settings and different environments. Nevertheless, our empirical results are already useful as a stepping stone for our second contribution, namely to identify *i*) the challenges and limitations that, once overcome, can potentially unlock performance gains beyond what reported here, and *ii*) preliminary techniques and ideas that can be exploited towards this end.

We argue that these two contributions can inspire a new and fruitful line of research aimed at further understanding, enhancing, and exploiting our concurrent ranging primitive.

2 Related Work

Our concurrent ranging primitive was originally inspired by the body of work on synchronous transmissions in narrowband low-power radios, pioneered by Glossy [4]. Synchronous transmissions exploit the PHY-level phenomena of constructive interference and capture effect to achieve unprecedented degrees of high reliability, low latency, and low energy consumption. Recent works [5, 6, 7] build atop the basic network flooding provided by Glossy to support more sophisticated traffic patterns.

A significant difference of this paper w.r.t. the literature on synchronous transmissions is that we concern ourselves with *ranging* rather than *communication*. This, combined with the focus on UWB radios and their fundamentally different principle of operation, yields a perspective far richer (and more challenging) than what hitherto explored. Indeed, the aforementioned existing works exploit synchronous transmissions as a way to boost packet reception rate (*PRR*), essentially regarding communication as a “black box” whose only relevant aspect is the outcome (i.e., success vs. failure). In this paper, by exploiting the availability of CIR information on the DW1000, we go beyond this perspective and consider concurrent packet exchanges as a “white box” whose innards can be analyzed and exploited to extract the timing information key to ranging. In doing so, we are in line with the current research trends exploiting

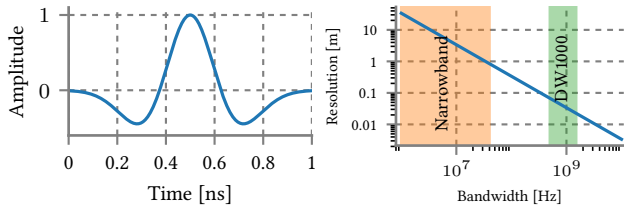


Figure 2. UWB pulse. **Figure 3. Distance resolution vs. bandwidth.**

channel state information (CSI) in narrowband technology for the purpose of localization [8, 9, 10].

This radically different perspective also sets us apart from the only related piece of work connecting synchronous transmissions and UWB, namely, the SurePoint system described in [11], which builds a Glossy-like flooding mechanism to *schedule* ranging exchanges, and exploits the capture effect for the data collection of timing measurements from anchors. In other words, SurePoint exploits synchronous transmissions as an efficient communication layer, as in the narrowband related literature. In contrast, and for the first time in the literature, we exploit synchronous transmissions to *perform* ranging exchanges, with the goal of removing altogether the need for scheduling, and the associated (and significant) latency and energy overhead.

3 UWB and the DecaWave DW1000

UWB communications have been originally used for military applications due to their very large bandwidth and interference resilience to mainstream narrowband radios. In 2002, the FCC approved the unlicensed use of UWB under strict power spectral masks, boosting a new wave of research from industry and academia. Nonetheless, this research mainly focused on high data rate communications, and remained largely based on theory and simulation, as most UWB radios available then were bulky, energy-hungry, and expensive, hindering the widespread adoption of UWB. In 2007, the IEEE 802.15.4a standard amendment included a UWB PHY layer based on impulse radio (IR-UWB), aimed at providing accurate ranging with low-power consumption. A few years ago, DecaWave released a standard-compliant IR-UWB radio, the DW1000, saving UWB from a decade-long oblivion, and taking to storm the field of real-time location systems (RTLS).

Impulse Radio. In IR-UWB, nodes transmit a time-hopping sequence of very short (≤ 2 ns) pulses (Figure 2) that translate into a large bandwidth (> 500 MHz) in the frequency domain [12]. The resulting large bandwidth leads to a very high time resolution (Figure 3), enabling UWB radios to measure a fine-grained version of the CIR. This, in turn, helps distinguish the signal’s leading path from multipath components, and therefore to accurately estimate the time-of-arrival (ToA) of a signal at the receiver. For instance, by analyzing the CIR, the DW1000 can estimate the ToA with a precision of 15 ps. Providing this time resolution usually entails incorporating high-speed ADCs that are costly and energy-hungry. Yet, the DW1000 can provide energy savings compared to the Intel 5300 (Table 1), the most representative WiFi radio able to measure CSI. For instance, to transmit a single packet

Table 1. Current consumption comparison. Note that the consumption depends on radio configuration.

State	DW1000	TI CC2650 [13]	Intel 5300 [14]
Deep Sleep	802.15.4a	BLE 4.2 & 802.15.4	802.11 a/b/g/n
Sleep	50 nA	100–150 nA	N/A
Idle	1 μ A	1 μ A	30.3 mA
TX	12–18 mA	550 μ A	248 mA
RX	35–85 mA	6.1–9.1 mA	387–636 mA
	57–126 mA	5.9–6.1 mA	248–484 mA

the DW1000 reduces the current consumption by a 85–91x factor, while providing an order of magnitude improvement in the ranging accuracy. On the other hand, communication in state-of-the-art low-power radios (e.g., the TI CC2650 in Table 1) is still significantly cheaper than with the DW1000.

Channel Impulse Response. The DW1000 measures the CIR upon reception of a signal, and stores it in a large internal buffer of 4064B. The measured CIR spans a symbol time, i.e., 993.59 ns for a 16 MHz pulse repetition frequency (PRF) or 1017.63 ns for a 64 MHz PRF [3]. This corresponds to 992 and 1016 samples respectively, where each sample is a complex number (16-bit integer real part and 16-bit imaginary part). The CIR measured by the DW1000 can be expressed as:

$$h = \sum_{k=1}^N a_k + jb_k = \sum_{k=1}^N A_k e^{j\theta_k} \quad (2)$$

where $A_k = \sqrt{a_k^2 + b_k^2}$ is the amplitude and $\theta_k = \arctan \frac{b_k}{a_k}$ is the phase of the signal; hereafter, we focus on the former. Upon reception of a packet, the DW1000 exploits the measured CIR information to properly estimate the leading path of the signal, which is exposed to the application as the first path index in the CIR buffer. Note that even when reception (RX) errors occur, the DW1000 measures the CIR, which can be therefore analyzed also in cases where a packet is *not* successfully decoded. However, in this case the DW1000 LDE microcode, which estimates the first path of the signal, does not compute the RX timestamp.

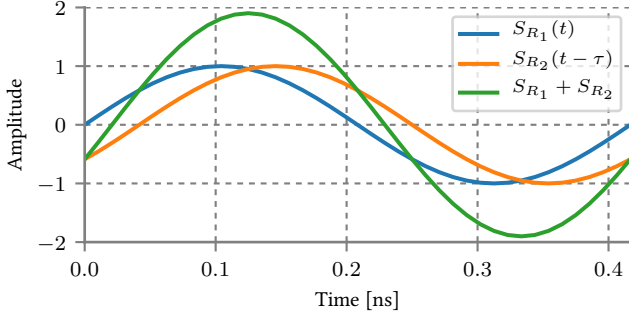
4 Concurrent Ranging

We describe next our concept of concurrent ranging and discuss its implementation in the DW1000.

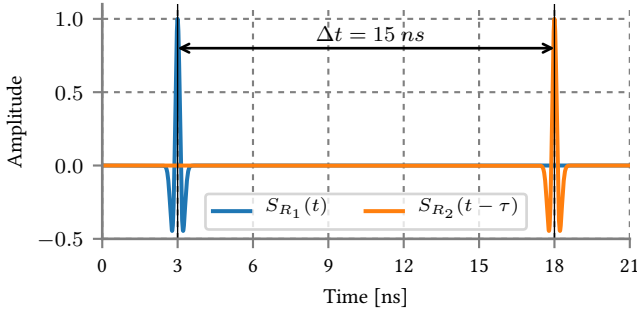
Concept. At its core, our idea of concurrent ranging is extremely simple. In a nutshell, it changes the basic single-sided two-way ranging (SS-TWR) scheme discussed in the introduction (Figure 1) by

1. replacing the N unicast POLL packets necessary to solicit ranging from the N responders with a single broadcast POLL packet, and
2. having all responders reply to the POLL packet after the same time interval δ_{TX} from its (timestamped) receipt.

It is worth noting that this simple idea is impractical (if not infeasible) in narrowband radios, as illustrated by the idealized view in Figure 4a, where the signals from responders R_1 and R_2 interfere with each other and are “fused” into a signal where the time information necessary to estimate distance is essentially lost. In contrast, Figure 4b shows that this is not the case in UWB; the time displacement Δt of the



(a) Narrowband: it is basically infeasible to recover the timing information of the individual responders' signals.



(b) UWB: the different distance from the initiator to responders R_1 and R_2 produces a time shift Δt between the transmissions of each responder. In UWB, this translates to two pulses separated by Δt . By measuring Δt , we can determine the distance difference $\Delta d = |d_1 - d_2|$ between responders.

Figure 4. Concurrent ranging, idealized.

two pulses, caused by their different distances from the initiator, is still clearly visible in the resulting signal, due to the fact that UWB pulses are extremely short w.r.t. narrowband waves, and therefore unlikely to interfere—although in practice things are not so simple, as discussed in Section 6.

As we discuss in detail in Section 6.3, the time shift Δt can be measured from the CIR as the difference between the position of the first pulse, directly obtained from the DW1000 along with the corresponding distance estimate d_1 , and the position of the second pulse, that we can estimate based on CIR information. Once Δt is known, the spatial displacement Δd can be computed, and the distance of the second responder obtained as $d_2 = d_1 + \Delta d$.

Implementation. Concurrent ranging can be implemented on the DW1000 in a very simple way, directly inherited by the simplicity of the concept, by replacing in the implementation of SS-TWR (e.g., the one provided by DecaWave) the unicast POLL packet with a broadcast one. As for the value of δ_{TX} , in our implementation we use $\delta_{TX} = 330 \mu s$, as we verified experimentally that this provides a good tradeoff; lower values would not leave enough time for the decoding of the POLL packet, and larger ones would negatively affect ranging due to clock drift and the need to retain nanosecond-level time synchronization.

Ideally, the relation $t_{R_{TX}} = t_{P_{RX}} + \delta_{TX}$ holds, with reference to our illustration of SS-TWR in Figure 1, yielding the time at which the RESPONSE packet must be sent. However, this time relies on two components. The first one is the timestamp $t_{P_{RX}}$ associated to the RX of the POLL packet

that, as mentioned in Section 3, the DW1000 can estimate with an extremely high precision of 15 ps. Unfortunately, the same precision is not provided for the TX of a packet, for which the DW1000 chip offers a scheduling granularity of 8 ns [1, 11]. In practice, this means that the actual timestamp at which the RESPONSE packet is transmitted is $t_{R_{TX}} \pm 8$ ns.

5 Open Questions

Although the idea of concurrent ranging is extremely simple and can be implemented straightforwardly on the DW1000, several questions must be answered to ascertain its practical feasibility. We discuss them next, before providing answers based on empirical observations.

Is communication even possible? Up to this point, we have implicitly assumed that the UWB transceiver is able to successfully decode one of the concurrent transmissions with high probability, similarly to what happens in narrowband and exploited by Glossy and other protocols. However, this is not necessarily the case, given the different radio PHY and the different degree of synchronization (nanoseconds vs. microseconds) involved.

How concurrent transmissions affect ranging accuracy? We also similarly implicitly assumed that concurrent transmissions do not affect the ranging accuracy. In practice, however, UWB pulses are far from being as “clean” as in the idealized view of Figure 4b. The main pulse is typically followed by several multipath reflections, which effectively create a “tail” after the main signal. Depending on its temporal and spatial displacement, this tail may interfere with the main pulse of other responders by *i*) reducing the size of the peak signal, and *ii*) generating peaks that can be mistaken for leading signals, inducing estimation errors.

Does the CIR contain enough information for ranging? Provided that concurrent transmissions do not yield a useless ranging accuracy, it is not for granted that the information in the CIR is of sufficient quality to recover the timing information necessary for distance estimation. In Section 4 we have mentioned that the limitation on the granularity of TX scheduling in the DW1000 introduces an 8 ns uncertainty. Given that an error of 1 ns in estimating the time of flight results in a corresponding error of ≈ 30 cm, this raises questions to whether concurrent ranging is actually feasible in practice.

Does the number of and relative distance among responders matter? If multiple responders are at a similar distance from the initiator, their pulses are likely to overlap in the CIR. This makes it difficult to discriminate among them to estimate the responders' distances, as well as to distinguish the leading signals from multipath components. Dually, if the distance between the initiator and the nearest responder is much smaller than the one w.r.t. other responders, the initiator may not discriminate their transmissions; the difference in power loss over distance may render the transmissions of the farther responders too faint to be detected, effectively masked by those of the nearest responder. The signals from responders that are much farther than the nearest one may even fall out of the CIR span.

We provide answers to these questions via the empirical observations in Section 6 confirming that, under rather permissive conditions, concurrent ranging can in principle yield

good accuracy. Nevertheless, these observations only confirm that the physical PHY-level phenomena are compatible with our proposed ranging primitive. A separate question is how to *design* a system that is actually capable of automatically and reliably extract the required timing information, i.e., *how to reliably detect the signal peaks associated to responders*. As this is the first paper investigating the notion of concurrent ranging, we do not yet have an exhaustive answer to this question. Nonetheless, we do offer some promising directions and preliminary techniques we discuss in Section 7.

6 Empirical Observations

We answer the open questions in Section 5, based on observations in a small-scale controlled experimental setting.

Hardware. All our experiments employ the DecaWave EVB1000 development platform [15], equipped with the DW1000 UWB transceiver and an STM32F105 ARM Cortex M3 microcontroller. We use the rectangular PCB antenna already present on the EVB1000 nodes.

UWB radio configuration. In all experiments, we use a preamble length of 128 symbols and a data rate of 6.8 Mbps. Further, we use channel 4, whose wider bandwidth² of 1.3 GHz (e.g., compared to the 500 MHz bandwidth of channel 1–3 and 5) provides better resolution in positioning the leading signal and therefore better ranging estimates.

Firmware. We program the behavior of initiator and responder nodes directly atop DecaWave libraries, without any OS layer, by adapting towards our goals the demo code provided by DecaWave. Specifically, we provide support to log, via the USB interface, the *i*) packets transmitted and received, *ii*) ranging measurements, and *iii*) the CIR measured upon reception of a packet.

Environment. All our experiments are carried out in a university building, and specifically in a long corridor whose width is 2.37 m. This is arguably a challenging environment due to the presence of strong multipath effects. Nevertheless, we also argue that it is a very realistic environment to test our ideas in, given that one of the main applications of UWB is for localization in GPS-denied environments, for which indoor in general and office buildings in particular constitute a prominent example.

Network configuration. In all the experiments we report about we have one initiator node and one or more responders. All the nodes are arranged in a line, placed exactly in the middle of the aforementioned corridor. Figure 5 shows for instance the network used in Section 6.1; we change the arrangement and number of nodes throughout the paper depending on the phenomenon under investigation. This one-dimensional configuration allows us to clearly and intuitively relate the temporal displacements of the received signals to the spatial displacement of their source nodes. We verified that this does not induce a bias w.r.t. a star topology retaining the same distances (e.g., d_1 and d_2 in Figure 5) used in the corresponding line configuration.

6.1 Is Communication Even Possible?

Our first goal is to ascertain whether UWB receivers are able to successfully decode a packet in the presence of con-

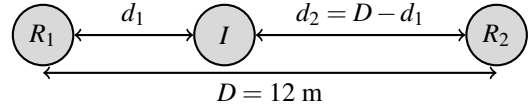


Figure 5. Experimental setup to investigate the reliability and accuracy of concurrent ranging (Section 6.1–6.2). I is the initiator, R_1 and R_2 are the responders.

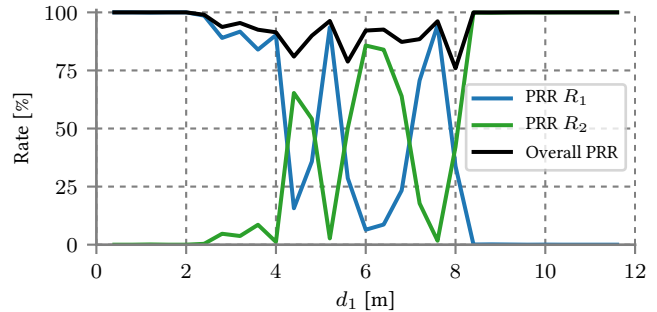


Figure 6. Reliability of concurrent transmissions.

current transmissions. We run a series of experiments with three nodes, one initiator I and two concurrent responders R_1 and R_2 , placed along a line (Figure 5). The initiator is placed between responders at a distance d_1 from R_1 and $d_2 = D - d_1$ from R_2 , where $D = 12$ m is the fixed distance between the responders. We vary d_1 between 0.4 m and 11.6 m in steps of 0.4 m. By changing the distance between the initiator and the responders we affect the chances of successfully receiving a packet from either responder due to the variation in power loss and propagation delay. For each initiator position, we perform 3000 ranging exchanges for statistical relevance, measuring the packet reception ratio (PRR) of the RESPONSE packets and the resulting ranging estimates. For all initiator positions, we performed 1000 ranging exchanges with each responder in isolation, always achieving $PRR = 100\%$ for all initiator positions.

Figure 6 shows the PRR of each responder and the overall $PRR_{R_1} + PRR_{R_2}$. Among all initiator positions, the worst overall $PRR = 75.93\%$ is achieved for $d_1 = 8$ m. On the other hand, when the initiator is close to one of the responders (i.e., $d_1 \leq 2$ m or $d_1 \geq 10$ m), the overall $PRR \geq 99.9\%$. We also observe strong fluctuations in PRR in the middle area. For instance, R_1 and R_2 achieve, respectively, a PRR of 93.6% and 2.7% for $d_1 = 5.2$ m, and of 6.43% and 85.73% for $d_1 = 6$ m.

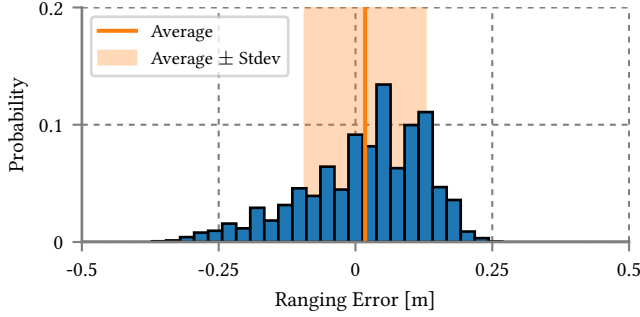
Summary. Overall, this experiment confirms the ability of the DW1000 to decode, with high probability, one of the packets from concurrent transmissions.

6.2 What about Ranging Accuracy?

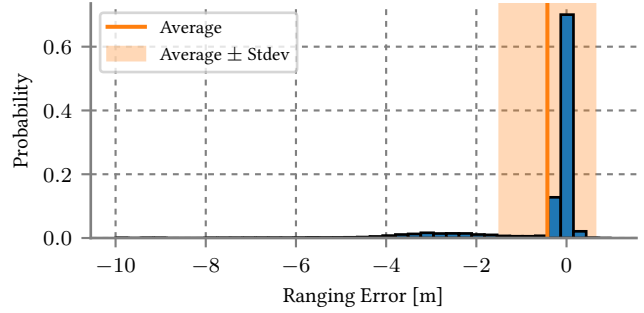
We now turn our attention to ascertaining whether concurrent transmissions degrade the ranging accuracy.

Baseline. We first look at the ranging accuracy for all initiator positions with each responder *in isolation*, using the same setup of Figure 5. Fig. 7a shows the normalized histogram of the resulting ranging error from 58000 ranging measurements. The average error is 1.7 cm, with a standard deviation

²Note that the DW1000 receiver bandwidth is limited to 900 MHz.



(a) Isolated responders.



(b) Concurrent responders.

Figure 7. Normalized histogram of the SS-TWR ranging error with responders in isolation (left) and two concurrent responders (right). When using concurrent responders, sometimes the initiator receives the packet from the farthest responder, while it estimates the first path for ranging from the closest responder, increasing the absolute error.

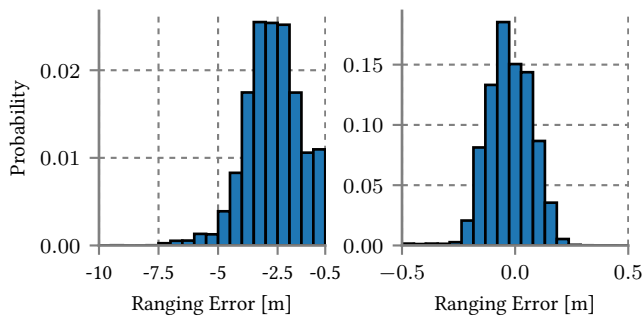


Figure 8. Zoomed-in views of Figure 7b.

of 10.9 cm. The maximum absolute error is 37 cm. The median of the absolute error is 8 cm, while the 99th percentile is 28 cm. These results are in accordance with previously reported studies [11, 16] employing the DW1000 transceiver.

Impact on ranging accuracy. Next, we look at the impact of concurrent transmissions on ranging. Fig. 7b shows the normalized histogram of the ranging error of 82519 measurements using two concurrent responders³. Overall, the median of the absolute error is 8 cm, as in the isolated case, while the 25th percentile and 75th percentiles are 4 cm and 15 cm, respectively. However, the average error is -0.42 cm and the standard deviation is 1.05 m. Further, the error distribution is clearly different w.r.t. the case of isolated responders (Figure 7a); to better appreciate the trends, Figure 8 offers a zoomed-in view of two key areas of the histogram in Figure 7b. Indeed, the latter has a long tail of measurements with significant errors; for 14.87% of the measured samples the ranging error is < -0.5 m, while in the isolated case the maximum absolute error only reaches 37 cm.

The culprit: Mismatch between received RESPONSE and nearest responder. To understand why, we study the ranging error when the initiator is located in the center area ($4 \leq d_1 \leq 8$), the one with major PRR fluctuations. Figure 9 shows the average absolute ranging error of the received packets from each responder as a function of the initiator position. Colored areas represent standard deviations.

The ranging error of R_1 and R_2 increases dramatically for $d_1 \geq 6$ m and $d_2 \geq 6$ m, respectively. Moreover, the magnitude of the error exhibits an interesting phenomenon. For instance, when the initiator is at $d_1 = 6.8$ m, the average error for the messages received from R_1 is 1.68 m, very close to the displacement between responders, $\Delta d = |d_1 - d_2| = |6.8 - 5.2| = 1.6$ m. Similarly, for $d_1 = 5.2$ m and $\Delta d = 1.6$ m, the average error for the messages received from R_2 is 1.47 m.

The observation that the ranging error approximates the displacement Δd between responders points to the fact that these high errors appear when the initiator receives the packet from the farthest responder but estimates the first path of the signal using the CIR peak corresponding instead to the nearest responder. This phenomenon explains the high errors shown in Fig. 7b and 8 (left), which are the result of this mismatch between the successful responder and the origin of the obtained first path. In fact, the higher probabilities in Fig. 8 (left) correspond to the positions where the farther responder from the initiator achieves the highest PRR in Figure 6. For example, for $d_1 = 7.6$ m, the far responder R_1 achieves $PRR = 94.46\%$ and an average ranging error of -3.27 m, which again corresponds to $\Delta d = 3.2$ m and also to the highest probability in Figure 8 (left).

The role of TX scheduling uncertainty. When this mismatch occurs, we also observe a relatively large standard deviation in the ranging error. This is likely generated by the 8 ns TX scheduling granularity of the DW1000 transceiver, discussed in Section 4. In SS-TWR (Figure 1), responders insert in the RESPONSE message the elapsed time δ_{TX} between receiving the POLL message and sending the response. The initiator uses δ_{TX} to precisely estimate the time-of-flight of the signal. However, the 8 ns uncertainty produces a discrepancy between the δ_{TX} used by the initiator and obtained from the successful response, and the δ_{TX} used by the closest responder, which results in significant error variations.

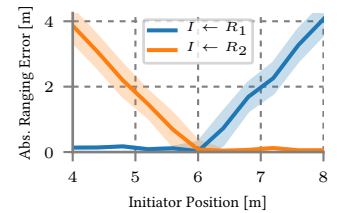


Figure 9. Ranging error vs. initiator position.

³Note we do not obtain valid ranging measurements in case of RX errors.

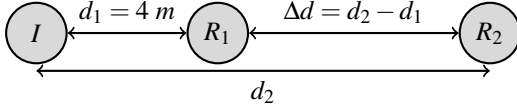


Figure 10. Experimental setup to analyze the CIR resulting from concurrent ranging (Section 6.3).

Summary. Concurrent transmissions can negatively affect ranging by producing a mismatch between the successful responder and the CIR peak used to compute the time of flight. However, we also note that 84.59% of the ranging samples are quite accurate, achieving an absolute error below 30 cm.

6.3 Is the Information in the CIR Enough?

To ascertain whether the CIR contains enough information to estimate the distance between each responder and the initiator, we run another series of experiments using again three nodes but deployed in a slightly different arrangement, depicted in Figure 10. In these experiments, we set the initiator I and the responder R_1 at a fixed distance $d_1 = 4$ m. We place the second responder R_2 at a distance $d_2 > d_1$ from the initiator; the two responders are therefore separated by a distance $\Delta d = d_2 - d_1$. Unlike previous experiments, we increase d_2 in steps of 0.8 m; we explore $4.8 \leq d_2 \leq 12$ m, and therefore $0.8 \leq \Delta d \leq 8$ m. For each position of R_2 , we run the experiment until we successfully receive 500 RESPONSE packets, yielding 500 valid ranging estimates; we measure the CIR on the initiator after every received response.

Baseline. Before using concurrent responders, we first measured the CIR of responder R_1 ($d_1 = 4$ m) in isolation. Figure 11 shows the average amplitude and standard deviation of 500 CIR signals; the signals are averaged by aligning them to the first path index reported by the DW1000. The measured CIR presents a clear leading path at 50 ns, followed by strong multipath components. We observe that the CIR barely changes across the 500 signals, exhibiting only minor variations in these multipath components (around 55–65 ns).

Concurrent responders: Distance estimation. We now analyze the effect of having a second responder R_2 transmitting *concurrently* with R_1 , and show how the corresponding distance can be estimated. We focus on a single distance $d_2 = 9.6$ m and on a single CIR (Figure 12), to analyze in depth the phenomena at stake; we later discuss results over 500 CIR signals (Figure 13) and for other d_2 values (Table 2).

Figure 12 shows that the response of R_2 introduces a second peak in the CIR, centered around 90 ns. This is compatible with our a-priori knowledge of $d_2 = 9.6$ m; the question is whether this distance can be estimated from the CIR.

Positioning the peak from R_2 in time constitutes a problem per se. In the case of R_1 , this estimation is performed accurately and automatically by the DW1000, which also returns the corresponding estimate of d_1 . The same could be performed for R_2 if it were in isolation, but not concurrently with R_1 . Therefore, we estimate the position of the peak from R_2 based on the CIR index whose signal amplitude is closest to 20% of the maximum amplitude of the peak—a widely-accepted technique used, e.g., in [17]. The offset between this CIR index and the one returned by the DW1000 for R_1 ,

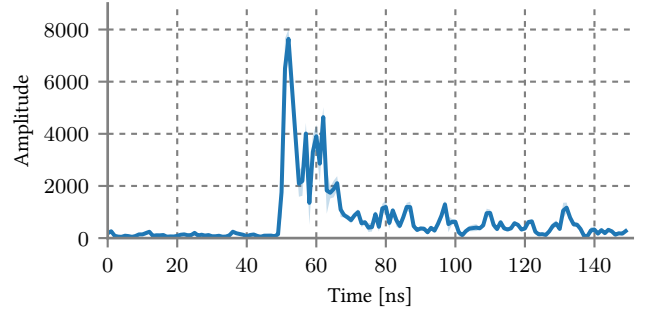


Figure 11. Average amplitude and standard deviation of 500 CIR signals for an isolated responder at $d_1 = 4$ m.

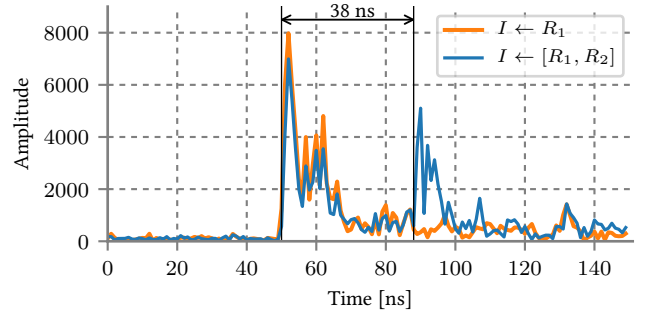


Figure 12. Impact of concurrent transmissions on the CIR.

for which a precise estimate is available, returns the delay Δt between the responses of R_1 and R_2 .

This delay is induced by the different propagation delay caused by the difference $\Delta d = d_2 - d_1$ in the distance of the responders from the initiator. Recall the basics of SS-TWR (Figure 1) and of our concurrent ranging primitive (Section 4). R_2 receives the POLL from I slightly after R_1 ; the propagation of the RESPONSE back to I incurs the same delay; therefore, the response from R_2 arrives at I with a delay $\Delta t = 2 \times \frac{\Delta d}{c}$ w.r.t. R_1 , where c is the speed of light in air.

In our case, the estimate above from the CIR signal yields $\Delta t = 38$ ns, corresponding to $\Delta d \approx 5.6$ m—indeed the displacement of the two responders. Therefore, knowing the distance d_1 from the initiator to R_1 , estimated precisely by the DW1000, we can easily estimate the distance between I and R_2 as $d_1 + \Delta d$. This confirms that a single concurrent ranging exchange contains enough information to reconstruct both distance estimates.

Concurrent responders: Sources of ranging error. Another way to look at Figure 12 is to compare it against Figure 4b; while the latter provides an *idealized* view of what happens on the UWB channel, Figure 12 provides a *real* view. The interference among the non-ideal pulses reduces the amplitude of the peak of R_1 ; it is therefore interesting to see whether this holds in general and also, dually, what is the impact on the (weaker) pulse of R_2 .

To this end, Figure 13 shows the average amplitude and standard deviation of 500 CIR signals with $d_1 = 4$ m, and $d_2 = 9.6$ m. We observe that the first pulse, the one from R_1 , presents only minor variations in the amplitude of the main

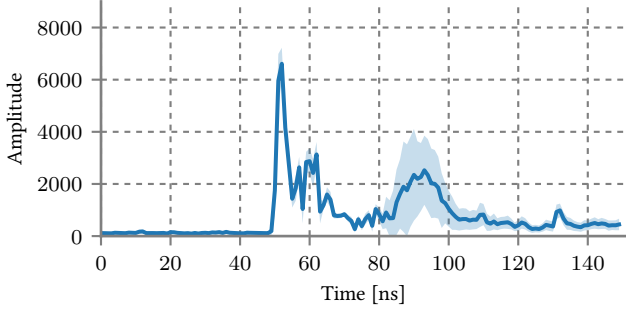


Figure 13. Average amplitude and standard deviation of 500 CIR signals for two concurrent responders at distance $d_1 = 4$ m and $d_2 = 9.6$ m from the initiator.

peak and of the multipath components, coherently with Figure 11. In contrast, the pulse from R_2 exhibits stronger variations, as shown by the colored area between 80 and 110 ns representing the standard deviation. However, these variations can only marginally be ascribed to interference with the pulse from R_1 ; we argue, and provide evidence next, that these variations are caused by the result of small time shifts of the observed CIR pulse, in turn caused by the 8 ns uncertainty in TX scheduling.

TX uncertainty affects time offsets. Figure 14 shows the normalized histogram, for the same 500 CIR signals, of the time offset Δt between the time at which the responses from R_1 and R_2 are received at I . The real value, computed by exact knowledge of distances, is $\Delta t = 37.37$ ns; the average from the CIR samples is instead $\Delta t = 36.11$ ns, with a standard deviation of 2.85 ns. These values, and the trends in Figure 14, are compatible with the 8 ns uncertainty deriving from TX scheduling.

Time offsets affect distance offsets. As shown in Figure 14, the uncertainty in time offset directly translates into uncertainty in the distance offset, whose real value is $\Delta d = 5.6$ m. In contrast, the average estimate is $\Delta d = 5.41$ m, with a standard deviation of 0.43 m. The average error is therefore -18 cm, with a median of 35 cm; the 75th and 99th percentile are 54 cm and 1.25 m, respectively. These results could still provide sub-meter ranging accuracy as long as the distance estimated to R_1 is accurate enough.

Distance offsets affect ranging error. Recall that the distance d_1 from R_1 to I is returned directly by the DW1000, while for R_2 is estimated as $d_2 = d_1 + \Delta d$. Therefore, the uncertainty in the distance offset Δd directly translates into an additional ranging error, shown in Figure 15 for each responder.

The plot shows that R_1 exhibits a ranging error of 3.6 cm with a standard deviation of 1.8 cm and a 99th percentile of 8 cm. As a result, the ranging error for R_2 achieves an average error of -15 cm with a standard deviation of 42.67 cm. The median error of R_2 is 31cm, while the 25th, 75th, and 99th percentiles are 16 cm, 58 cm, and 1.18 m.

Impact of distance between responders. In principle, the results above demonstrate the feasibility of concurrent ranging and its ability to achieve sub-meter ranging estimates, although less precise than conventional, isolated ranging.

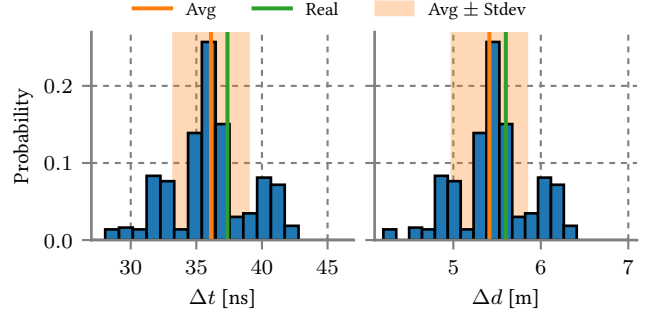


Figure 14. Normalized histograms of the time offset Δt and corresponding distance offset Δd between the leading CIR pulses from R_1 and R_2 .

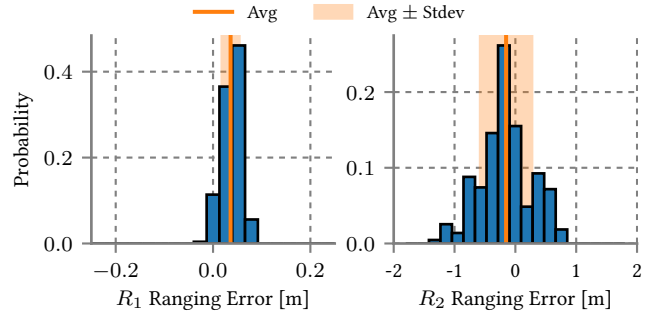


Figure 15. Normalized histograms of the concurrent ranging error of each responder.

Nevertheless, these results were obtained for a single value of d_2 . Table 2 summarizes the results we obtained by varying this distance as described at the beginning of the section. We only consider the RESPONSE messages successfully sent by R_1 , since those received from R_2 produce the mismatch mentioned in Section 6.2, increasing the error by $\approx \Delta d$. We discuss in Section 7.2 a preliminary technique that could be used to detect these mismatches and estimate the distance to each responder reliably.

To automatically compute the peaks and the leading paths, we exploit our a-priori knowledge of where the peaks should be located based on Δd , and therefore Δt . For each responder, we consider the slice of the CIR defined by $\Delta t \pm 8$ ns, and detect the first peak in it, estimating the leading edge as the preceding index with the amplitude closest to the 20% of the maximum amplitude, as described earlier. To abate false positives, we also enforce the additional constraints that a peak should have a minimum amplitude of 1500, and that the minimum distance between peaks is 8 ns.

The estimated distance to R_1 achieves an average error below 9 cm with a standard deviation below 10 cm for all tested d_2 distances. The 99th percentile absolute error is always below 27 cm. These results are in line with the ranging performance obtained in Section 6.2. We observe that the largest error of the estimated Δd and d_2 is obtained for the shortest distance $d_2 = 4.8$ m. In this particular setting, the CIR peaks corresponding to both responders are very close together and may even overlap, suffering as well the impact of strong multipath components and increasing the resulting

Table 2. Concurrent ranging performance.

d_2	Δd	PRR [%]			Estimated Δd [m]		R_1 Ranging Error [cm]					R_2 Ranging Error [cm]				
		PRR_{R_1}	PRR_{R_2}	Overall	Avg	Std	Avg	Std	50th	75th	99th	Avg	Std	50th	75th	99th
4.8	0.8	2.54	95.31	97.85	0.33	0.27	3	9	6	7	26	-43	32	30	73	105
5.6	1.6	36.3	36.73	73.03	1.5	0.38	6	2	6	8	12	-4	38	31	43	83
6.4	2.4	65.04	22.09	87.13	2.09	0.76	6	2	5	7	10	-25	76	51	113	161
7.2	3.2	0.2	99.6	99.8	3.0	0.0	8	0	8	8	8	-12	0	12	12	12
8.0	4.0	38.12	44.55	82.67	4.07	0.46	8	2	9	10	13	16	46	41	58	96
8.8	4.8	69.23	20.39	89.62	4.78	0.38	5	2	5	6	9	3	38	25	43	86
9.6	5.6	100.0	0.0	100.0	5.41	0.43	4	2	4	5	8	-15	43	31	58	118
10.4	6.4	94.76	2.52	97.28	6.42	0.44	5	2	5	7	9	7	44	36	53	99
11.2	7.2	85.05	5.23	90.27	7.16	0.4	6	2	6	7	10	2	39	34	42	97
12.0	8.0	100.0	0.0	100.0	8.06	0.35	4	2	5	6	9	11	35	29	44	77

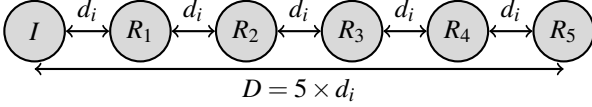


Figure 16. Experimental setup to analyze the CIR resulting from five concurrent responders (Section 6.4).

error (-43 cm in average for d_2). The rest of the distances achieve an average error below 26 cm. We observe that the error is significantly lower with $\Delta d \geq 4$ m, achieving a 75th percentile lower than 60 cm. Similarly, for all $\Delta d \geq 4$ m but $\Delta d = 5.6$ m, the 99th percentile is lower than 1 m. These results confirm that concurrent ranging can, indeed, achieve sub-meter ranging accuracy, at least if the distance Δd between the responders is sufficiently large.

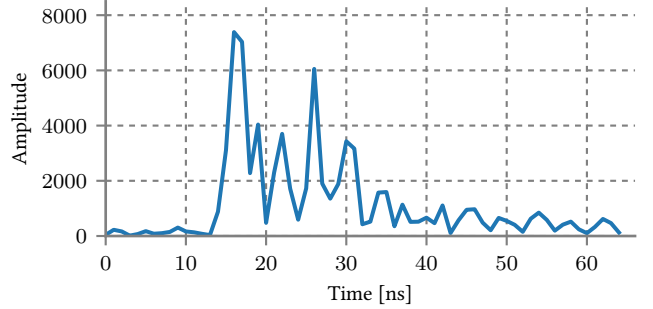
Summary. Concurrent ranging can achieve sub-meter accuracy, but requires *i*) a sufficiently large Δd between the concurrent responders, otherwise, the responders’ pulses may not be distinguishable, and *ii*) obtain the RESPONSE packet from the closest responder, otherwise, a mismatch can increase the ranging error to $\approx \Delta d$.

6.4 What about More Responders?

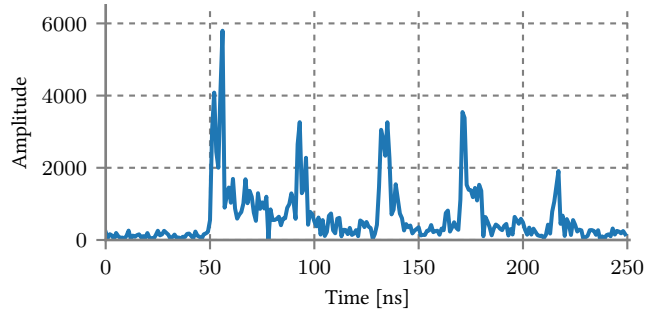
We conclude our experimental campaign by investigating the impact of more than two concurrent responders, and their relative distance, on PRR and ranging performance. To this end, we run experiments with five concurrent responders deployed in a line as depicted in Figure 16, for which we change the inter-node distance d_i . For every tested distance d_i , we repeat the experiment until we obtain 500 successfully received RESPONSE packets.

Dense configuration. We begin by examining a very short distance $d_i = 0.4$ m, yielding rather similar distances between each responder and the initiator. Recall that 1 ns ≈ 30 cm. Given that a UWB pulse is 1–2 ns, this means that the pulses from neighboring responders are likely to overlap with this choice of d_i .

The corresponding CIR in Figure 17a shows that this is indeed the case. Although we can visually observe different peaks, discriminating the ones associated to responders from those caused by multipath is very difficult if not impossible in absence of a-priori knowledge about the number of concurrent responders and/or the environment characteristics. Even when this is present, in some cases the CIR shows a wider pulse that “fuses” the pulses of one or more respon-



(a) $d_i = 0.4$ m: the peaks corresponding to each responder are not clearly distinguishable, distance from the initiator cannot be estimated.



(b) $d_i = 6$ m: the peaks corresponding to each responder are clearly separated, distance from the initiator can be estimated.

Figure 17. Impact of the relative distance d_i among 5 responders, analyzed via the corresponding CIR.

ders along with multipath components. In essence, when the difference in distance $\Delta d = d_i$ among multiple responders is too small, concurrent ranging cannot be applied. Interestingly, however, for $d_i = 0.4$ m the overall PRR = 99.36%.

Sparser configurations: PRR. We now explore $2 \leq d_i \leq 10$ m, therefore up to a maximum distance $D = 50$ m from the initiator to the last responder. The experiment achieved an overall PRR = 96.59%. The minimum PRR was 88.2% for the maximum $d_i = 10$ m, while the maximum PRR = 100% was achieved for $d_i = 8$ m. The closest responder R_1 achieved an overall PRR = 90.56%. The PRR of the experiment is interestingly high, especially considering that in narrowband technologies increasing the number of transmitters sending different packets typically decreases

Table 3. Concurrent ranging performance with five responders.

d_i	R_1 Error [cm]					R_2 Error [cm]					R_3 Error [cm]					R_4 Error [cm]					R_5 Error [cm]				
	Avg	Std	50th	75th	99th	Avg	Std	50th	75th	99th	Avg	Std	50th	75th	99th	Avg	Std	50th	75th	99th	Avg	Std	50th	75th	99th
2	-12	2	12	13	17	-42	49	51	76	164	-96	81	111	173	206	-85	67	86	147	207	-117	65	125	176	216
4	-3	2	2	4	8	-1	32	26	37	100	-21	39	25	52	110	-32	37	34	51	124	-39	32	33	59	118
6	2	2	1	4	7	4	37	23	46	100	-9	37	26	44	89	-14	44	30	54	121	-14	42	30	48	110
8	3	2	3	5	8	-23	41	28	58	122	5	40	32	51	94	-22	31	29	43	90	6	34	25	39	75
10	3	1	2	5	7	-15	42	23	52	107	-23	44	30	54	137	-14	33	27	44	96	-25	32	23	43	97

the performance due to scalability problems of the capture effect [5]. We associate this high PRR to the closer distance to the initiator of R_1 w.r.t. the other responders.

Sparsers configurations: Ranging error. Table 3 shows the ranging error for all positions and responders. We use the same technique as in Section 6.3 to compute the peaks and leading paths and, similarly, only consider the exchanges (about 90% in this case) where the successfully received RESPONSE packet is from the nearest responder R_1 , to remove the bias induced by the mismatch discussed in Section 6.2.

We observe that $d_i = 2$ m provides the worst performance; this is the result of the peaks from different responders being still quite close from each other, and therefore suffering from the multipath components of previously transmitted pulses.

On the other hand, Figure 17b shows an example CIR for $d_i = 6$ m, an intermediate value between the extremes considered. The CIR clearly shows five different peaks, which enables the initiator to measure the distance to each responder. We observe that the time offset between two consecutive peaks is similar; this is expected given the same distance offset $\Delta d = d_i$ between two neighboring responders.

In general, Table 3 shows that when $d_i > 2$ m, the average error for all responders remains below 40 cm, and the 75th percentile absolute error remains below 60 cm.

Summary. We argue that these results make sub-meter concurrent ranging a reality that can be exploited to balance energy consumption with accuracy. However, to be of practical use, we first need to design more sophisticated mechanisms that allow us to efficiently and reliably detect peaks and distinguish them from multipath components.

7 Reliably Detecting Responder Peaks

In this paper, we are primarily concerned with accruing empirical evidence confirming the feasibility of concurrent ranging, and enabling us to assess its limitations in terms of performance and applicability.

Nonetheless, exploiting concurrent ranging in *practical* systems entails building a toolchain that, based on CIR information, is able to *dynamically, automatically, and reliably* detect the peaks associated to responders, therefore enabling computation of their corresponding distance estimates.

In Section 7.1–7.2 we illustrate two techniques, and preliminary results, we consider a stepping stone towards this goal. Note that we are *not* proposing that these techniques fully solve the problem in isolation; we conjecture that the aforementioned toolchain will likely employ synergically several techniques. In this respect, we elaborate on other complementary techniques in Section 7.3.

7.1 Power Boundary

One of the key challenges to overcome lies in discriminating between the main CIR peaks associated to the leading path of concurrent signals and their multipath components.

We define a *power boundary* that allows us to discriminate the two; pulses that are the result of the leading path will likely show a signal amplitude above the boundary, while multipath components will likely remain below it.

This boundary can be defined, e.g., based on the widely-known Friis transmission equation that relates the power transmitted P_t and received P_r to the distance d between transmitter and receiver:

$$\frac{P_r}{P_t} = G_t G_r \left(\frac{\lambda}{4\pi d} \right)^2$$

Barring manufacturing issues and configuring all devices with the same settings, we can assume P_t and the other parameters λ (wavelength), G_t and G_r (antenna gains) to be the same for all transmissions. This allows us to relate the expected received power P_{r_1} from responder R_1 with the power P_{r_2} received from R_2 as a function of the distances d_1 and d_2 travelled by each signal, respectively, by dividing P_{r_1} by P_{r_2} :

$$\frac{P_{r_1}}{P_{r_2}} = \frac{d_2^2}{d_1^2}$$

Assuming d_1 is the distance of the nearest responder, the above can be rewritten into

$$P_{r_2} = P_{r_1} \frac{d_1^2}{d_2^2} = P_{r_1} \frac{d_1^2}{(d_1 + \Delta d)^2} \quad (3)$$

by defining $d_2 = d_1 + \Delta d$, as we discussed in Section 6.3.

Eq. (3) provides the desired power boundary by relating the signal decay of the main peak of R_1 with the position (i.e., distance, and therefore time displacement) in the CIR of the signals following this peak. Multipath components remain below this boundary. In contrast, observe that while the two leading paths from concurrent transmissions incur a time offset $\Delta t = 2 \times \frac{\Delta d}{c}$, the distance difference travelled by their signals is only Δd . As a result, peaks from concurrent transmissions should have a relatively high power for their position (time displacement) in the CIR, likely placing them above the boundary and therefore enabling their discrimination from multipath components.

Figure 18 illustrates the concept on a real CIR signal obtained with responders R_1 and R_2 at $d_1 = 4$ m and $d_2 = 8$ m from the initiator. The plot shows the power boundary computed from Eq. (3), based on the distance d_1 and the power P_{r_1} , both made available by the DW1000 when estimating the

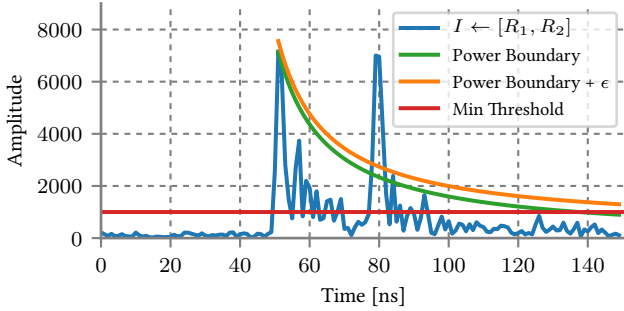


Figure 18. Power boundary to distinguish CIR peaks introduced by concurrent transmissions from multipath.

distance of the first responder R_1 . The peak from R_2 (around 80 ns) is well above the power boundary, as well as the minimum power threshold we introduce to avoid the many false positives multipath would induce over long distances.

On the other hand, the Friis equation is ideal, and may not fully capture the peculiarity of the environment at hand. In this respect, we envision two routes to increase the robustness in filtering out multipath components. On one hand, the boundary can be extended upwards (also shown in Figure 18) by an environment-dependent additive component ϵ , whose value should be calibrated empirically in the target environment. Alternately, the Friis equation could be replaced by the path loss equation [18]; the latter is commonly used in indoor localization [19], but nonetheless requires a similar calibration of environment-dependent parameters. Commonly used RF models [20] could be also used to identify and discard clusters of multipath components.

7.2 Signal Cross-correlation

Concurrent ranging can provide sub-meter ranging accuracy, but cannot associate the identity of a responder to a distance measurement. To tackle this issue, we propose to exploit the signal cross-correlation between the acquired CIR signal with concurrent responders and a previously-obtained CIR for each isolated responder. These CIRs in isolation could be obtained, e.g., by application traffic or dedicated ranging exchanges using the conventional SS-TWR.

In signal processing, cross-correlation is typically used to find where two signals match or present the maximum similarity. Therefore, when the CIR for an isolated signal from a responder R_i is matched against the CIR of several concurrent responders including R_i , it should exhibit a maximum where the two signals are more similar, enabling the association of R_i with a peak in the concurrent CIR.

Figure 19 exemplifies the concept by showing the cross-correlation of the two CIR signals in Figure 12, viz. *i*) one from an isolated responder R_1 at $d_1 = 4$ m from the initiator, and *ii*) the other with two concurrent responders at $d_1 = 4$ m and $d_2 = 9.6$ m. The cross-correlation presents an absolute maximum in correspondence of a time shift between the two signals of 1 ns. This maximum is the result of the very high similarity between the CIR pulse with R_1 in isolation and the first pulse in the CIR with two concurrent responders, which corresponds to R_1 , and directly inform us that the first peak in the CIR with concurrent responders corresponds to R_1 .

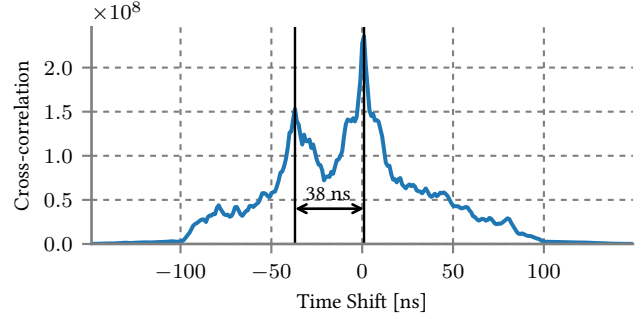


Figure 19. Cross-correlation of the two CIR signals presented in Figure 12. The maximum peak directly identifies the first responder when using concurrent transmissions.

Further, cross-correlation shows a relative maximum with a time shift of -37 ns, a result of the relatively high cross-correlation between the peak introduced by R_2 and the CIR with R_1 in isolation. We also observe that the time shift between the two maxima is exactly 38 ns, i.e., the theoretical Δt determined by the distances d_1 and d_2 above. As a result, cross-correlation could also help the initiator to improve the estimate of Δd between the different responders.

Finally, we argue that cross-correlation could be also exploited to detect the mismatch introduced in Section 6.2, which makes many ranging exchanges error-prone. Using cross-correlation, we can identify the CIR peak that corresponds to the successful responder, i.e., the one from which the RESPONSE packet was actually received. This would enable the initiator to remove the error introduced by the mismatch, and measure reliably the distance to each responder.

7.3 Other Complementary Approaches

We briefly outline other approaches that can be used in conjunction with the techniques we just described.

Interleaving isolated and concurrent ranging. As mentioned in Section 7.2, isolated ranging can be used to acquire a CIR and perform cross-correlation. However, isolated ranging can also be useful as a fallback option in situations where low *PRR* or high fluctuations in ranging accuracy are detected. Further, it can be deliberately interleaved with concurrent ranging, effectively using the former to limit the error induced by the latter, yet achieving energy savings w.r.t. the conventional case.

Exploiting knowledge about the environment. If the absolute position of the responders is known to the initiator, this information can be used when computing its relative distance to them, and rule out ranging estimates that would be impossible to satisfy in the environment at hand. Other environmental information, e.g., the placement of walls as in our experimental setup, could also be exploited to estimate the temporal displacement of multipath components via geometrical considerations [21]. This could be used to filter out multipath signals and improve the detection of the main responder peaks.

8 Discussion

The results we present here must be confirmed by additional experiments in other environments, with different topology configurations, and possibly alternate radio configurations. However, they clearly show that concurrent ranging strikes a novel and potentially disruptive tradeoff between ranging accuracy and energy consumption of UWB radios. As mentioned in Section 1, the potential benefits are significant, as they entail reducing energy consumption and latency by a factor of N , being the number of target responders.

On the other hand, the results we present here also show that the benefits of concurrent ranging are not immediately available. A number of challenges must be overcome to improve the reliability of the primitive; for some of them we have proposed preliminary techniques in Section 7. We contend that the full exploitation of the potential of concurrent ranging is likely to be enabled by a synergistic application of several of these techniques.

The evolution of UWB transceivers may also partially remove or simplify some of these challenges. For instance, reducing the 8 ns uncertainty associated to TX scheduling in the DW1000 may already significantly improve the accuracy of concurrent ranging, as discussed in Section 6.3. Similarly, enhancing the ns-level time resolution provided in the CIR towards the ps-level resolution internally used by the DW1000 estimation would greatly improve the accurate detection of the leading path of concurrent responders beyond the first.

Finally, and most importantly, we do *not* necessarily consider concurrent ranging as a *replacement* of conventional ranging. There will always be applications for which the best accuracy is necessary. Nevertheless, we contend that the tradeoffs unlocked by concurrent ranging empower designers with extra degrees of freedom to trade accuracy vs. energy, enabling new strategies for reconciling these conflicting concerns in ways hitherto impossible.

9 Conclusions

We presented a novel concurrent ranging primitive for UWB radios that holds the potential to redefine the tradeoffs between accuracy, latency, and energy consumption. The concept underlying this primitive and its implementation is simple, replacing the pairwise, unicast exchanges of the popular single-sided two-way ranging scheme with a broadcast POLL triggering a concurrent RESPONSE from neighbors.

Our empirical observations from a small-scale experimental setup confirm that concurrent ranging can achieve sub-meter accuracy. However, our results are also the basis to elicit a few key challenges that hamper the immediate exploitation of this primitive, for which we proposed some techniques and preliminary results.

10 References

- [1] DecaWave. DW1000 Data Sheet, 2016.
- [2] IEEE Std 802.15.4-2011. *IEEE Standard for Local and metropolitan area networks—Part 15.4: Low-Rate Wireless Personal Area Networks (LR-WPANs)*, 2011.
- [3] DecaWave. DW1000 User Manual, 2016.
- [4] F. Ferrari, M. Zimmerling, L. Thiele, and O. Saukh. Efficient Network Flooding and Time Synchronization with Glossy. In *Proc. of the 10th ACM/IEEE Int. Conference on Information Processing in Sensor Networks (IPSN)*, 2011.
- [5] O. Landsiedel, F. Ferrari, and M. Zimmerling. Chaos: Versatile and Efficient All-to-all Data Sharing and In-network Processing at Scale. In *Proc. of the 11th ACM Conference on Embedded Networked Sensor Systems (SenSys)*, 2013.
- [6] F. Ferrari, M. Zimmerling, L. Mottola, and L. Thiele. Low-power Wireless Bus. In *Proc. of the 10th ACM Conference on Embedded Network Sensor Systems (SenSys)*, 2012.
- [7] T. Istomin, A. L. Murphy, G. P. Picco, and U. Raza. Data Prediction + Synchronous Transmissions = Ultra-low Power Wireless Sensor Networks. In *Proc. of the 14th ACM Conference on Embedded Network Sensor Systems (SenSys)*, 2016.
- [8] S. Sen, B. Radunovic, R. R. Choudhury, and T. Minka. You Are Facing the Mona Lisa: Spot Localization Using PHY Layer Information. In *Proc. of the 10th Int. Conference on Mobile Systems, Applications, and Services (MobiSys)*, 2012.
- [9] D. Vasisht, S. Kumar, and D. Katabi. Decimeter-level Localization with a Single WiFi Access Point. In *Proc. of the 13th Usenix Conference on Networked Systems Design and Implementation (NSDI)*, 2016.
- [10] M. Kotaru, K. Joshi, D. Bharadia, and S. Katti. SpotFi: Decimeter Level Localization Using WiFi. In *Proc. of the ACM Conference on Special Interest Group on Data Communication (SIGCOMM)*, 2015.
- [11] B. Kempke, P. Pannuto, B. Campbell, and P. Dutta. SurePoint: Exploiting Ultra Wideband Flooding and Diversity to Provide Robust, Scalable, High-Fidelity Indoor Localization. In *Proc. of the 14th ACM Conference on Embedded Network Sensor Systems (SenSys)*, 2016.
- [12] M. Z. Win and R. A. Scholtz. Impulse radio: How it works. *IEEE Communications letters*, 2(2):36–38, 1998.
- [13] Texas Instruments. CC2650 SimpleLink™ Multistandard Wireless MCU, 2016.
- [14] Halperin, D. and Greenstein, B. and Sheth, A. and Wetherall, D. Demystifying 802.11n Power Consumption. In *Proc. of the Int. Conference on Power Aware Computing and Systems (HotPower)*, 2010.
- [15] DecaWave. DecaWave ScenSor EVB1000 Evaluation Board, 2013.
- [16] B. Kempke, P. Pannuto, and P. Dutta. PolyPoint: Guiding Indoor Quadrotors with Ultra-Wideband Localization. In *Proceedings of the 2nd Int. Workshop on Hot Topics in Wireless (HotWireless)*, 2015.
- [17] B. Kempke, P. Pannuto, and P. Dutta. Harmonium: Asymmetric, Bandstitched UWB for Fast, Accurate, and Robust Indoor Localization. In *Proc. of the 15th Int. Conference on Information Processing in Sensor Networks (IPSN)*, 2016.
- [18] S. Y. Seidel and R. S. Rappaport. 914 MHz path loss prediction models for indoor wireless communications in multifloored buildings. *IEEE transactions on Antennas and Propagation*, 40(2):207–217, 1992.
- [19] N. Patwari, J. N. Ash, S. Kyperountas, A. O. Hero, R. L. Moses, and N. S. Correal. Locating the nodes: cooperative localization in wireless sensor networks. *IEEE Signal processing magazine*, 22(4):54–69, 2005.
- [20] A. A. Saleh and R. Valenzuela. A Statistical Model for Indoor Multipath Propagation. *IEEE Journal on selected areas in communications*, 5(2):128–137, 1987.
- [21] E. Leitinger, M. Fröhle, P. Meissner, and K. Witrisal. Multipath-assisted maximum-likelihood indoor positioning using UWB signals. In *Proc. of the IEEE Int. Conference on Communications Workshops (ICC)*, 2014.