Paper published at the 38th IEEE International Conference on Distributed Computing Systems (ICDCS). July 2–5, 2018.

Concurrent Ranging with Ultra-Wideband Radios: From Experimental Evidence to a Practical Solution

Bernhard Großwindhager, Carlo Alberto Boano, Michael Rath, and Kay Römer
Faculty of Electrical and Information Engineering, Graz University of Technology, Austria
E-mail: {grosswindhager, cboano, mrath, roemer}@tugraz.at

***Abstract*—To enable future location-aware Internet of Things (IoT) applications, Ultra-wideband (UWB) technology provides centimeter-accurate distance estimations. In the common case of a non-synchronized network, at least $N \cdot (N-1)$ message exchanges are required to derive the distance between N nodes. Enabling concurrent ranging between an initiator and an arbitrary number of responders can drastically reduce the amount of necessary transmissions and hence increases the efficiency of UWB systems. Although the feasibility of concurrent ranging has been proven experimentally, several key challenges still need to be addressed to practically implement concurrent ranging in real-world UWB systems, such as the automatic detection of multiple responses, the identification of a responder, as well as the detection of overlapping responses (especially in the presence of multipath components). In this paper, we provide a concurrent ranging solution tackling the aforementioned challenges. Among others, our solution enables (i) to detect responses in the CIR reliably, (ii) to encode the responder ID in the CIR to allow personalized ranging, as well as (iii) to mitigate the impact of overlapping responses and multipath components. We further show how the proposed solution increases the scalability of concurrent ranging in real-world UWB-based distributed systems.**

***Index Terms*—Channel impulse response, concurrent ranging, Decawave DW1000, Multipath components, Ultra-wideband.**

I. Introduction

Ultra-wideband (UWB) radios allow for precise ranging and localization thanks to the high bandwidth ($\geq$ 500 MHz). Theoretical work on UWB technology dates back to the late 1990s and early 2000s [1], [2]. However, UWB was not commercially successful [3] until (i) the release of the IEEE 802.15.4a amendment adding a UWB physical layer [4] and (ii) the commercialization of the first low-cost IEEE 802.15.4-compliant UWB transceiver, the Decawave DW1000 [5]. These two key drivers have allowed UWB to become a key technology enabling location-aware IoT applications with centimeter-level positioning accuracy [6].

Scheduled ranging. The distance between two UWB nodes in a non-synchronized network is estimated by carrying out a two-way ranging scheme, i.e., a pair-wise exchange of at least two messages between two nodes: an initiator and a responder. To estimate the distance between all N nodes in a network, one typically needs to schedule the exchange of $N \cdot (N-1)$ messages. This procedure is not only time-consuming; it is especially energy-inefficient, considering that the DW1000 radio draws up to 155mA and 90mA in receive and transmit mode, respectively. This is significantly higher than for other low-power wireless technologies such as BLE and LoRa [7].

Concurrent ranging. The short pulses transmitted by UWB radios reduce multipath fading and, hence, allow to resolve individual multipath components. Recent work has used this capability to build multipath-assisted indoor localization systems [8], [9]. This property can also be used to extract the simultaneous responses of an arbitrary number of responders. This principle is called *concurrent ranging* and its feasibility was shown experimentally by Corbalán and Picco [10]. Instead of scheduling several ranging operations between an initiator and other responders in the network, the latter reply simultaneously to a single broadcast message. The different responses (i.e., the transmitted pulses of the preamble from each responder) are visible in the channel impulse response (CIR) of the initiator. Hence, by detecting these pulses (i.e., by identifying the signal peaks associated to the responders in the channel impulse response estimated by off-the-shelf UWB transceivers), it is possible to estimate the path delay and distance to all responders concurrently.

Open challenges. Although the feasibility of concurrent ranging was shown experimentally by Corbalán and Picco [10], several key challenges still need to be addressed to practically implement concurrent ranging in real-world UWB systems.

I. Automatic detection of multiple responses. To practically implement concurrent ranging, it is necessary to automatically detect the responses of several nodes in the received CIR. In other words, an initiator should be able to process the estimated CIR *at run-time* and reliably identify the signal peaks associated to the different responders. Although existing work has discussed possible approaches to achieve this goal [10], a practical implementation working at run-time is still missing.

II. Identifying responders. A key challenge hindering the feasibility of concurrent ranging in real-world systems is the impossibility of associating a distance estimate to a specific responder, i.e., the anonymity of ranging. Previous work investigating concurrent ranging applied artificial setups where all nodes are placed in a line [10], which gives the initiator prior knowledge about the order in which the signal peaks associated to the responders are received in the CIR. In practical situations, however, the relative locations of nodes are typically unknown. Approaches making use of cross-correlation between the CIR acquired with concurrent responders and a previously-obtained CIR for each isolated responder [10] are also not applicable, as the channel impulse response for an isolated responder varies depending on its position and on the surrounding environment.

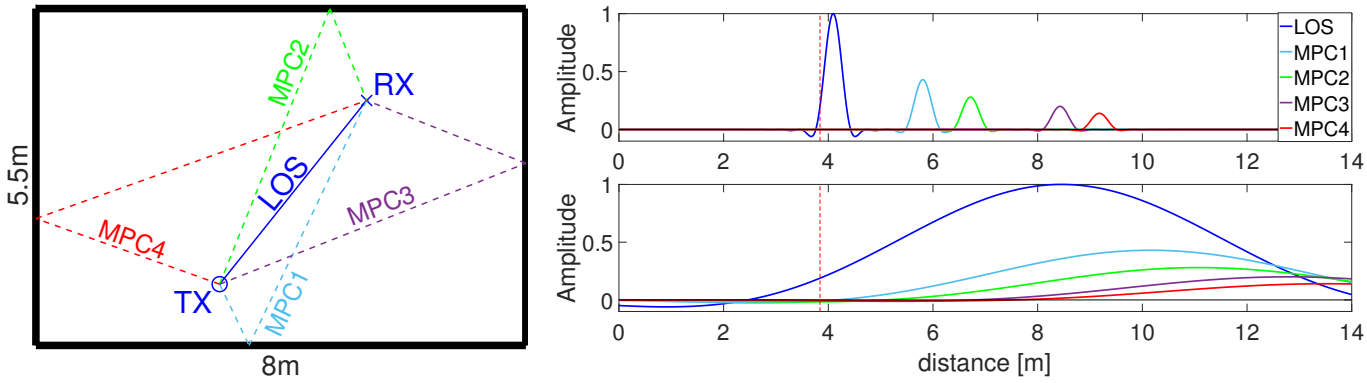

(a) Floor plan

(b) Theoretical multipath reflections

Fig. 1: (a) Floor plan showing line-of-sight (LOS) path and first-order reflections (MPC1 – MPC4); (b) Theoretical multipath reflections with a bandwidth of 900 MHz (top) and 50 MHz (bottom).

III. Detecting overlapping responses. If several responders are at a similar distance from the initiator, their responses will overlap, making it difficult to extract meaningful information from the channel impulse response. To date, no practical solution to decode overlapping responses has been presented.

IV. Mitigating the impact of multipath reflections. Another open challenge is how to differentiate between a response and a strong multipath component from another responder. Corbalán et al. [10] suggested to use power boundaries based on the Friis equation to differentiate between main CIR peaks and disturbing multipath components. This principle has three main issues limiting its applicability in real-world networks: (i) the Friis equation is idealized and does not hold true in typical UWB operational areas; (ii) in the case of an attenuated direct path due to obstacles, it is likely that multipath reflections have higher amplitudes than the direct path; (iii) the amplitude of the peaks in a CIR derived from low-cost UWB transceivers is highly varying. Hence, there is a need for algorithms that operate independently of the absolute amplitude of the signal peaks associated to the different responders.

V. Scalable concurrent ranging. Existing work has not yet focused on how to maximize the number of responders performing concurrent ranging. In current solutions, an initiator can obtain responses from multiple responders, but only as long as the latter are physically far away from each other, which limits the applicability of concurrent ranging in typical indoor UWB application settings.

Contributions. In this paper, we tackle all aforementioned challenges and provide a solution that allows the practical implementation of concurrent ranging on off-the-shelf UWB devices. We hence significantly advance the state-of-the-art in UWB concurrent ranging by addressing the yet open challenges highlighted by recent feasibility studies [10]. First, we describe a method to let initiators automatically detect responses in the channel impulse response at run-time *independently* of their absolute amplitude, hence making concurrent ranging feasible in real-world UWB applications (Sect. IV). Second, we present a novel technique based on *pulse shaping* that allows to associate a distance measurement to a specific responder, so that ranging is no longer anonymous (Sect. V). Third, we show that the employed algorithm allows to detect the signal peaks associated to the different responders reliably even in the case of overlapping responses (Sect. VI). Fourth, we propose a method to prevent the overlap of responses and strong multipath components from other responders by employing response position modulation (Sect. VII). Finally, we show that, by combining response position modulation and pulse shaping, we can provide a *scalable* concurrent ranging solution that can be practically implemented on off-the-shelf UWB devices (Sect. VIII). Before presenting our contributions in detail, we provide the reader with basic information about ultra-wideband technology (Sect. II) and concurrent ranging (Sect. III).

II. Ultra-Wideband Basics

The use of a high bandwidth (and consequently very short pulses) in UWB transceivers allows to resolve individual multipath components (MPC) as illustrated in Figure 1. In particular, Figure 1a shows a rectangular floor plan with a transmitter (TX) sending pulse signals to a receiver (RX). Due to the omni-directional wave propagation and to the reflections from walls, multiple versions of the same pulse arrive at the receiver. Figure 1a shows the line-of-sight (solid) as well as the first-order MPCs (dashed), i.e., the pulse is reflected only at a single object. Figure 1b shows the (theoretically) received pulses. The top figure shows pulses with a bandwidth of 900 MHz (i.e., the maximum bandwidth of the DW1000 radio [5]), whilst the bottom figure shows pulses sent with a bandwidth of 50 MHz. Due to the steep rising edge at 900 MHz, the precision of the distance estimation in a system based on time-of-flight (ToF) is increased.

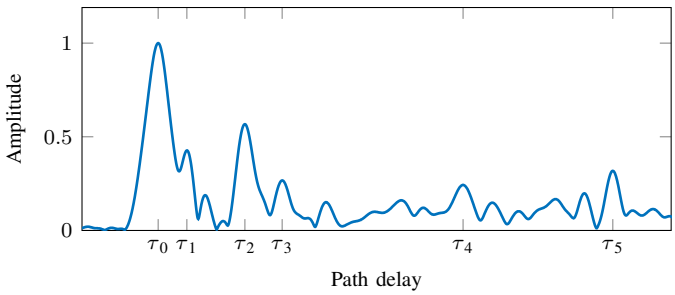

Fig. 2: Estimated CIR obtained from the DW1000 radio in an indoor environment. The path delay of multipath components is marked with τ_k, where τ_0 denotes the LOS component.

Furthermore, the multipath reflections at 50 MHz are highly overlapping. This makes narrowband systems more susceptible to multipath fading and – in contrast to UWB systems – it is not possible to extract multipath components. This is an important observation, as widely-used low-power wireless technologies such as BLE have a bandwidth that is even smaller than 5 MHz.

At a reasonably high bandwidth, Figure 1b resembles the channel impulse response (CIR), i.e., information about the multipath propagation consisting of reflections from walls and scattering from other objects. The channel impulse response $h(t)$ can be modeled as [8]:

$$h(t) = \sum_{k=0}^{K} \alpha_k \delta(t - \tau_k) + \nu(t) \quad (1)$$

with α_k and τ_k denoting, respectively, the complex-valued amplitude and path delay of K deterministic multipath components resulting from specular reflections from walls, windows, or doors. The last term $\nu(t)$ is diffuse or non-deterministic multipath, i.e., higher-order reflections or signal components due to scattering. UWB transceivers such as the Decawave DW1000 provide a channel impulse response estimation to precisely estimate the arrival time of a packet by detecting the first pulse in the CIR. The DW1000 radio estimates the arrival time with a precision of 15.65ps (using a 63.9 GHz sampling clock), which results in a distance resolution of 4.69mm [11]. Figure 2 shows an exemplary estimated channel impulse response obtained with the DW1000 radio. It depicts the line-of-sight (LOS) component (τ_0) and significant multipath reflections ($\tau_1 - \tau_5$). Besides estimating the arrival time of a signal, the CIR information can be used to enable multipath-assisted UWB localization [8], [9]. Furthermore, the CIR can be used to detect a degrading channel as well as any change of the surrounding environment: this can be exploited to adapt UWB physical layer (PHY) parameters and increase communication performance [7]. In this paper, instead, we use the CIR to receive simultaneous responses from several neighbors in a single packet and to estimate the distance of a wireless node to each of the neighbors simultaneously, as shown in Sect. III.

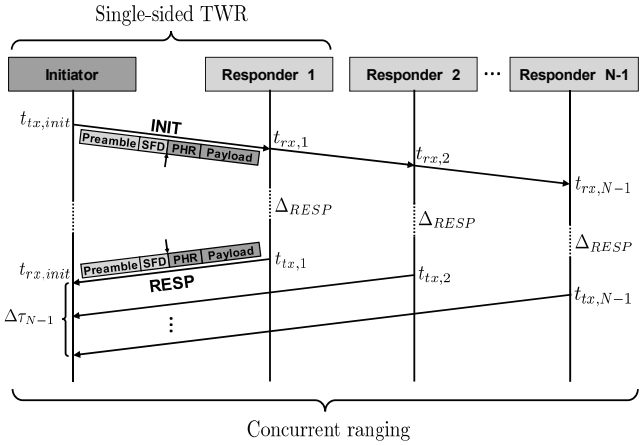

Fig. 3: Principle of single-sided two-way ranging (SS-TWR) and concurrent ranging schemes.

III. Concurrent Ranging

The single-sided two-way ranging (SS-TWR) scheme traditionally used to derive the distance between two users in a non-synchronized network requires the exchange of two messages (*INIT, RESP*) between an initiator and a responder (see Figure 3). Thus, in a network with N users, each node requires $N - 1$ transmissions and receptions.

Given the high current draw of UWB transceivers (especially in receive mode [7]), it is important to reduce the number of messages exchanged in a network to make UWB systems feasible for building location-aware IoT applications. Furthermore, scheduling the distance estimation to each neighbor requires a significant amount of time, which increases channel utilization and traffic load, as well as the inaccuracy of ranging in mobile settings.

In a *concurrent ranging* scheme, instead, the initiator broadcasts the *INIT* message to all neighbors (responders), who reply simultaneously with a *RESP* message after a constant delay Δ_{RESP}. Consequently, the *RESP* messages (containing the timestamps $t_{rx,i}$ and $t_{tx,i}$ in the payload) sent by all responders are overlapping in time. In narrowband transceivers, this leads to (unusable) severely overlapping pulses, as shown in Figure 1b. Using ultra-wideband radios, instead, it is possible to identify the signal peaks associated to each responder in the estimated channel impulse response. By employing concurrent ranging, the total number of messages needed to estimate distances to all neighbors is hence reduced from $N \cdot (N-1)$ to N. In fact, the initiator has to broadcast just one message and, more importantly, to receive just a single message that aggregates all responses.

Figure 3 also shows the UWB PHY frame structure according to the IEEE 802.15.4 standard. It consists of the preamble, start-of-frame delimiter (SFD), physical layer header (PHR), as well as a payload. The channel impulse response, used to derive responses from multiple neighbors, is estimated solely from the preamble and is hence independent from the payload.

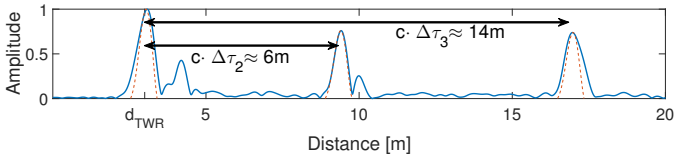

(a) Channel impulse response with three responders separated by 3m each. The dashed orange line marks fitted pulse templates.

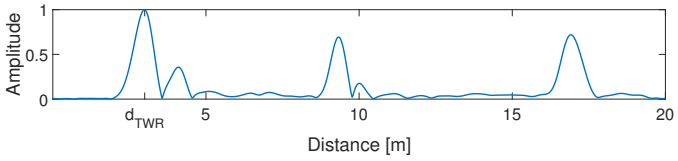

(b) Matched Filter output.

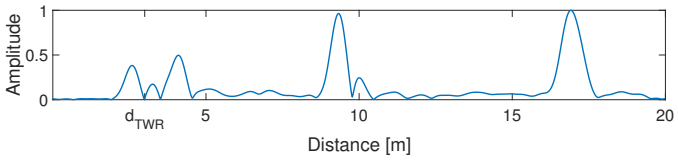

(c) Matched Filter output after subtracting strongest peak.

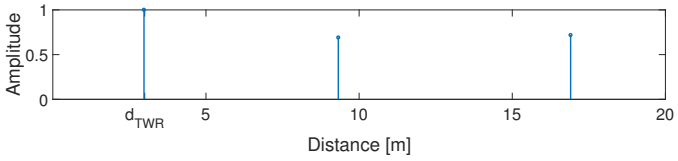

(d) Detected Responses.

Fig. 4: Principle of the proposed response detection algorithm.

Thus, the extracted responses from the CIR correspond to the pulses transmitted in the preamble (consisting of a predefined sequence of single pulses). The payload, instead, is not sent as single pulses, but as bursts of pulses [7].

The IEEE 802.15.4 standard defines the timestamp of a frame as the beginning of the first symbol of the PHR (RMARKER, marked with arrow in Figure 3) [12]. Hence, the minimum delay Δ_{RESP} consists of the duration of PHR and payload of *INIT* message, as well as the preamble and SFD of the *RESP* message. Using a data rate of $DR = 6.8$ Mbps, a pulse repetition frequency $PRF = 64$ MHz, and a preamble symbol repetition $PSR = 128$, this results in a duration of $178.5\mu s$. Additionally, we evaluated experimentally the minimum time necessary to switch the DW1000 radio from receive to transmit mode, which is less than $100\mu s$. Including a safety gap, we hence set the delay Δ_{RESP} to 290μs.

Figure 4a shows an acquired (normalized) channel impulse response from the *RESP* message when three neighbors are responding concurrently in a hallway. The three responders are placed at a distance from the initiator of $d_1 = 3m$, $d_2 = 6m$, and $d_3 = 10m$, respectively. Three significant peaks are visible in the CIR shown in Figure 4a, representing the strongest path component of each neighbor. The distance between the initiator and responder 1 is derived from the SS-TWR scheme (see Figure 3), as it is still possible to decode one of the concurrently transmitted payloads containing the required timestamps [10]. The formula to calculate the distance between initiator and responder 1 is:

$$d_{TWR} = \frac{(t_{rx,init} - t_{tx,init}) - (t_{tx,1} - t_{rx,1})}{2} \cdot c \qquad (2)$$

with c denoting the propagation speed in air. However, the distance between the initiator and the remaining responders is derived from the CIR. Due to the longer time-of-flight, the responding peaks of responder 2 and 3 arrive with a delay of $\Delta\tau_2 = 2 \cdot (\tau_2 - \tau_1)$ and $\Delta\tau_3 = 2 \cdot (\tau_3 - \tau_1)$ at the initiator, with τ_i ($i \in \{1, \ldots, 3\}$) denoting the path delay between the intiator and each responder. The resulting delays $\Delta\tau_2$ and $\Delta\tau_3$ are due to the *INIT* as well as the *RESP* message: hence, they have to be halved to correctly estimate the distance. The resulting distance between initiator and responders is thus $d_2 = d_{TWR} + \frac{c \cdot \Delta\tau_2}{2} = 6m$ and $d_3 = d_{TWR} + \frac{c \cdot \Delta\tau_3}{2} = 10m$, respectively.

Limited TX timestamp resolution. The Decawave DW1000 UWB transceiver has the useful feature of *delayed transmission*. The latter enables to set a future timestamp at which the transceiver sends a packet. This allows to align a pre-calculated timestamp with the real transmit timestamp and embed it in the message being transmitted ($t_{tx,i}$ in Figure 3). Unfortunately, the Decawave DW1000 ignores the low-order 9 bits of the delay transmit value, limiting the transmission timestamp resolution to approximately 8ns [11, p. 26]. This is not an issue in the classical SS-TWR scheme as the real transmit timestamp is anyway embedded in the message, but it has a severe impact on the precision of the concurrent ranging scheme, as it negatively affects the concurrency of the *RESP* messages of the neighbors. Given that this issue is hardware-dependent and could be solved in the next-generation UWB transceivers, this problem is out of scope of this paper.

IV. Reliable Response Detection

To make concurrent ranging feasible in real-world systems, it is essential to detect responses from the neighbors reliably in the channel impulse response. To this end, we propose a scheme based on the *search and subtract* algorithm [13]. This algorithm employs a matched filter computing the correlation between the received CIR and a pulse template with duration T_p transmitted by the UWB radio. To detect the $N-1$ strongest responses in the CIR, we use the following procedure:

1) We first upsample the CIR using *fast Fourier transform* in order to obtain a smoother signal. Furthermore, to correct for the unknown time offset of the CIR derived from the DW1000 radio [8], the channel impulse response is aligned with the distance d_{TWR} (see Eq. 2 and Figure 4). This step is not necessarily required, as the differences of the peaks are relevant for concurrent ranging, but it enhances visualization of the CIR and simplifies plausibility checks of the result.

2) Denoting the pulse template as $s(t)$, we can define the time-discrete impulse response of the matched filter as the time-reversed pulse template:

$$\mathbf{h_{MF}} = [s((N_p - 1) \cdot T_s), s((N_p - 2) \cdot T_s), ..., s(0 \cdot T_s)]$$

with T_s marking the sampling period, and $N_p = T_p/T_s$ the number of samples of the pulse. The output of the matched filter $\mathbf{y}$ is computed as the discrete convolution ($*$) between the impulse response of the matched filter $\mathbf{h}_{MF}$ and the derived CIR denoted as $\mathbf{r}$

$$\mathbf{y} = \mathbf{h}_{MF} * \mathbf{r}. \tag{3}$$

Figure 4b shows the matched filter output of the CIR depicted in Figure 4a. It is evident that the matched filter increases the signal-to-noise ratio (SNR) of the channel impulse response.

3) We identify the sample corresponding to the maximum of the matched filter output $\mathbf{y}$, indicating the index of the strongest path l_k. The latter relates to the path delay with $\tau_k = l_k \cdot T_s$ and to the path length with $d_k = \tau_k \cdot c$.
4) To reduce complexity, instead of the least squares solution suggested in [13], we calculate the amplitude of the strongest path $\hat{\alpha}_k$ as the amplitude of $\mathbf{y}$ at sample l_k.
5) The estimated neighbor response ($\hat{\alpha}_k s(t - \tau_k)$) is subtracted from the received signal $\mathbf{r}$. Figure 4c shows the matched filter output of the remaining signal after subtracting the strongest peak/neighbor response.
6) We repeat steps 2 to 5 with the remaining output signal until the $N - 1$ strongest paths are detected. Figure 4d shows the $N - 1 = 3$ strongest peaks corresponding to the responses from the three neighbors.
7) Independently of their amplitude $\hat{\alpha}_k$, the responses ($\hat{\alpha}_k$,τ_k) are arranged in ascending order starting with the one of the closest neighbor. Being $\hat{\alpha}_1$ and τ_1 the amplitude and path delay of responder 1, respectively, the distance of responder i is calculated as:

$$d_i = d_{TWR} + \frac{c \cdot (\tau_i - \tau_1)}{2}. \tag{4}$$

Due to lack of information from Decawave regarding the transmitted pulse used in the DW1000 radio, we identified the pulse shape $s(t)$ used in step 2 of our detection algorithm with a measurement campaign. In particular, we connected transmitter and receiver with a SMA cable and a 60 dB attenuator to avoid reflections and saturation of the transceiver, respectively. We then let the receiver log 1000 CIRs and, in a post-processing step, cut the direct path component from the channel impulse response, and calculate the average pulse shape. Figure 5a shows the default pulse shape at Channel 7 (900 MHz bandwidth).

Following the algorithm described in this section, we are hence able to reliably detect the responses of all neighbors in the CIR to perform concurrent ranging.

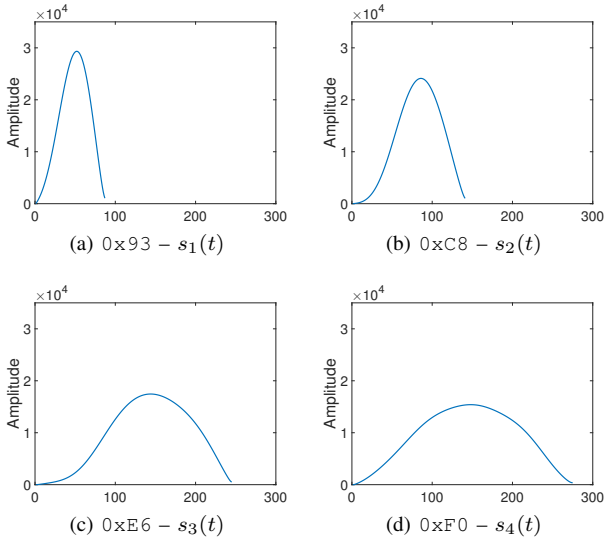

(a) `0x93` – $s_1(t)$ (b) `0xC8` – $s_2(t)$

(c) `0xE6` – $s_3(t)$ (d) `0xF0` – $s_4(t)$

Fig. 5: Pulse shape $s_i(t)$ for different values of the `TC_PGDELAY` register. `0x93` represents the default value.

V. Encoding Responder ID in the CIR

To make concurrent ranging usable in real-world applications, it is required to encode the ID of the responder in the channel impulse response. The preamble, indeed, consists of a fixed sequence of pulses (see Sect. III) and does not embed the identity of the sender. Hence, in the concurrent ranging scheme, the responses of the neighbors derived from the CIR do not contain any information to associate the responses to the corresponding neighbors. Consequently, distance estimations are typically anonymous in a concurrent ranging scheme. To remedy this problem, we suggest to use *pulse shaping* to associate a peak in the channel impulse response to a responder, i.e., to change the shape of the transmitted pulses as a function of the responder ID.

Pulse shaping. The off-the-shelf DW1000 radio provides the ability to change the width of the transmitted pulses via the 8-bit register `TC_PGDELAY`. Changing the value of this register effectively alters the output bandwidth [11, p. 148]. While making the pulse narrower (i.e., increasing the bandwidth) is not an option due to the regulatory spectral mask, making the pulse wider, instead, does not violate the regulations. Figure 5 exemplarily shows the pulse shape[1] $s_i(t)$ obtained when using the same settings as in Sect. IV (i.e., Channel 7, $PRF = 64$ MHz), and when configuring `TC_PGDELAY` with values `0x93` (s_1), `0xC8` (s_2), `0xE6` (s_3), and `0xF0` (s_4). Given that the default value (`0x93`) is the lower limit, up to 108 different pulse shapes are supported, which limits the theoretical number of possible responders differentiable with the proposed pulse shaping technique.

[1]Please note that (i) the amplitude of the pulses shown in Figure 5 are different due to scaling the pulses to unit energy, and that (ii) `0x93` is the default value of the `TC_PGDELAY` register for the employed configuration.

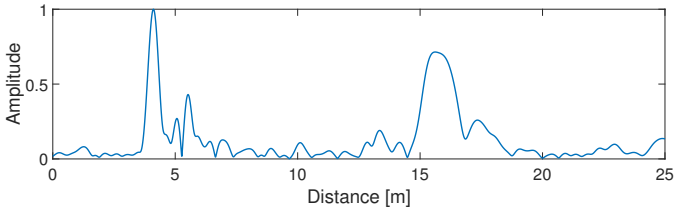

(a) CIR with two responders at $d_1 = 4m$ and $d_2 = 10m$, respectively.

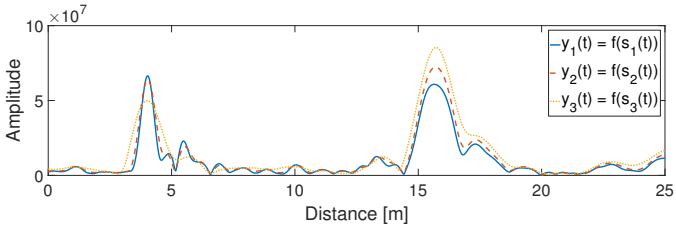

(b) Matched Filter output with $s_1(t)$, $s_2(t)$, and $s_3(t)$.

Fig. 6: CIR and matched filter output in the case of two responders replying with different pulse shapes.

No impact on ranging performance. Increasing the default value of `TC_PGDELAY` results in a wider pulse, which reduces the bandwidth and hence, in theory, also the ranging precision. To evaluate if the change in pulse shape has any impact on ranging performance, we place two UWB nodes three meters apart from each other in an office environment, and perform 5000 SS-TWR operations with three different pulse shapes (s_1,s_2,s_3 in Figure 5). Our results show that the standard deviation of the difference between the true distance and the estimated distance for the three pulse shapes is $\sigma_1 = 0.0228m$ (s_1), $\sigma_2 = 0.0221m$ (s_2), and $\sigma_3 = 0.0283$ (s_3), respectively. Therefore, changing the pulse shape by configuring `TC_PGDELAY` has a negligible impact on the ranging precision and can be safely used to encode the responder ID in a concurrent ranging scheme.

Identifying pulse shapes. Figure 6a shows the acquired channel impulse response when a neighbor at a distance $d_1 = 4m$ answers using the default pulse $s_1(t)$ (`0x93`, see Figure 5a), and a second neighbor at a distance of $d_2 = 10m$ responds with a wider pulse $s_3(t)$ (`0xE6`, see Figure 5c). The different pulse shape is clearly visible in the channel impulse response. Performing the algorithm described in Sect. IV with $N_{PS} = 3$ possible pulse templates $s_i(t)$ (with $i \in \{1, \ldots, N_{PS}\}$) results in the matched filter outputs $y_i(t)$ shown in Figure 6b. In each case, the responses of the neighbors are easily detectable, independently of the pulse template. To decode the transmitted pulse shape of the responders and hence their ID, we compare the estimated amplitudes of the neighbor responses $\hat{\alpha}_{k,i}$ (with k denoting the number of the response) of all N_{PS} matched filter outputs $y_i(t)$. The pulse shape i at which the amplitude $\hat{\alpha}_{k,i}$ is the maximum determines the pulse shape used by the responder. Therefore, in Figure 6b, the first response corresponds to the pulse template $s_1(t)$ (blue, solid) and the second to the pulse template $s_3(t)$ (yellow, dotted). Please note that, in the provided example, the number of pulse templates is set to $N_{PS} = 3$, but, in principle, up to 108 concurrent responders can be supported.

$\mathbf{d_2}$ [m]	6	7	8	9	10
$\mathbf{s_2(t)}$ (`0xC8`) [%]	99.9	99.5	99.8	100	99.8
$\mathbf{s_3(t)}$ (`0xE6`) [%]	99.2	99.7	99.9	100	100

TABLE I: Percentage of pulse shapes identified correctly.

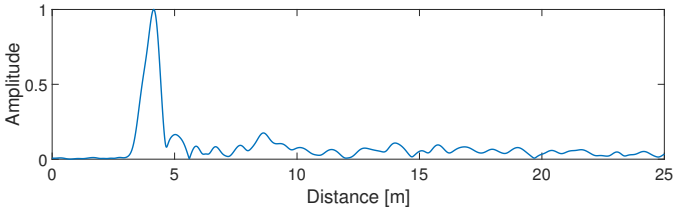

(a) Matched Filter output when two responses are overlapping.

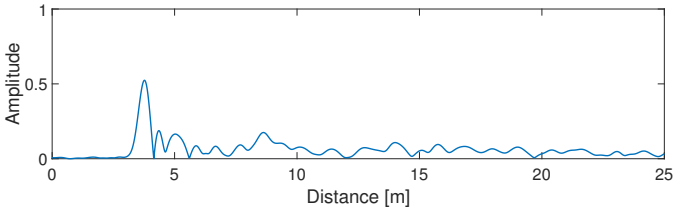

(b) Matched Filter output after subtracting strongest response.

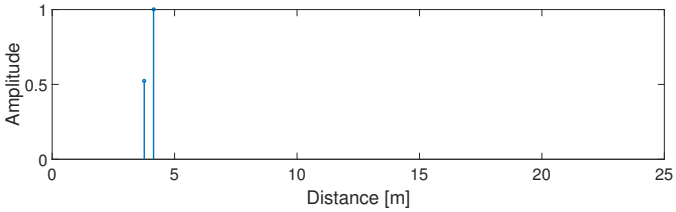

(c) Detected responses, both responses are detected successfully.

Fig. 7: Performance of the proposed algorithm in the case of overlapping responses from multiple responders.

Evaluation. We evaluate the performance of the proposed technique by placing an initiator and a responder at a fixed distance $d_1 = 3m$. Another responder is placed at a variable distance $d_2 \in \{6, 7, 8, 9, 10\}m$ from the initiator. Responder 1 uses the default pulse shape $s_1(t)$ (see Figure 5a). Responder 2 uses either $s_2(t)$ (see Figure 5b) or $s_3(t)$ (see Figure 5c). For each distance and pulse shape, we perform 1000 concurrent ranging operations. Table I shows how many pulses could successfully be identified. Regardless of the pulse shape and of the distance, a responder could successfully be identified in more than 99.2% of the cases, showing that *pulse shaping* can effectively encode the identity information of a responder in a concurrent ranging scheme.

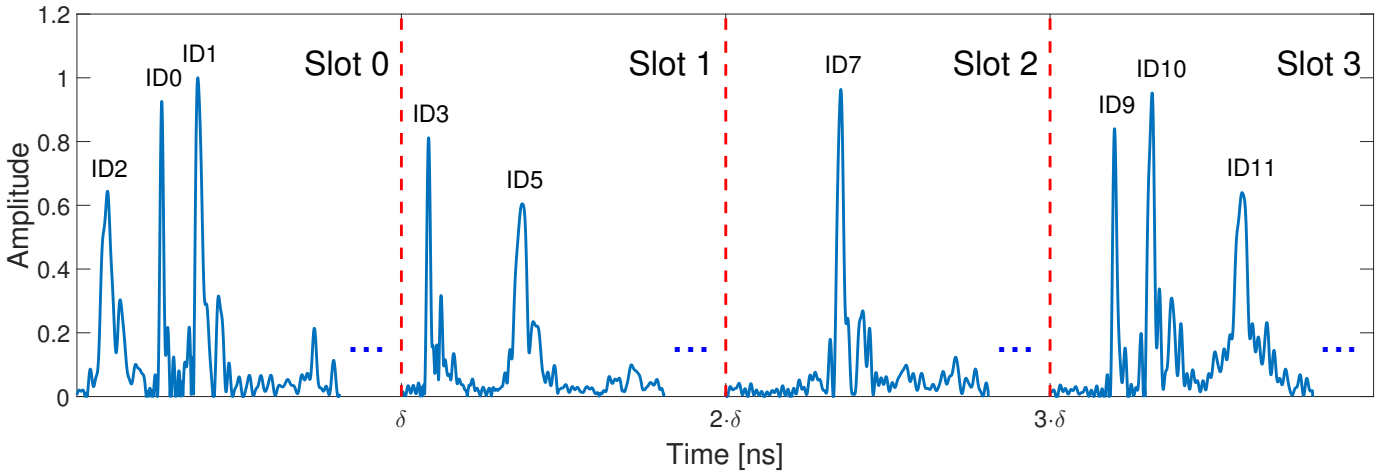

Fig. 8: Combining response position modulation with pulse shaping. In this example, nine users perform concurrent ranging by employing *four* slots and *three* different pulse shapes. The responders with ID 0, 1, 2 make use of pulse shape $s_1(t)$, $s_2(t)$, and $s_3(t)$, respectively.

VI. Detection of overlapping responses

So far, we assumed that the responses are nicely separated in time, i.e., that responses from responders placed at similar distances from the initiator do not overlap. We study next the performance of the detection algorithm proposed in Sect. IV in the presence of overlapping responses and show that it reliably detects responses from nodes at similar distances.

In particular, we compare the performance of our proposed algorithm with a *threshold-based algorithm* as proposed by Falsi et al. [13]. The threshold-based algorithm compares the channel impulse response with a defined threshold. If the CIR crosses this threshold, the maximum of the following N_p samples, i.e., the pulse duration, is derived. This operation is repeated until $N-1$ peaks are detected.

Both the *search and subtract* and the *threshold-based* algorithm exhibit a good performance when the responses are well-separated from each other. However, as soon as the responses are overlapping due to a similar distance (and hence time-of-flight) of several responders, the algorithm proposed in Sect. IV outperforms threshold-based algorithms.

We show this by acquiring 2000 concurrent ranges from two responders placed at the same distance $d_1 = d_2 = 4m$ from the initiator. The responses of the two nodes are highly overlapping: as a result, only one peak is visible in Figure 7a. Figure 7b shows that, after subtracting the strongest response with our proposed algorithm (step 5 described in Sect. IV), we obtain an easily detectable second response. Our evaluation shows that the threshold-based algorithm detects both responses in only 48% of the trials, whilst our algorithm is successful in 92.6% of the cases. Therefore, the algorithm that we propose in Sect. IV can be also used to detect the presence of responders positioned at similar distances from the initiator.

Please note that, even if the nodes are physically positioned at the same distance, the responses may still not overlap consistently. The reason is the limited transmission timestamp resolution of the Decawave DW1000 discussed in Sect. III, i.e., even if the distance is the same and two responders reply simultaneously, there might be an offset of $\pm$8ns. For this reason, in our performance evaluation, we have considered only trials in which the responses are actually overlapping.

VII. Mitigating the Impact of strong Multipath

In multipath-rich environments it is likely that several strong multipath components are received. In fact, it might be even the case that a received MPC is stronger than the direct path component due to blocked and attenuated line-of-sight. In these situations, it can hence be challenging to differentiate between a response from a neighbor and a dominant MPC.

Response position modulation. The most effective solution in these scenarios is to avoid that multipath components are overlapping with responses from other responders by separating them in time. Hence, we propose *response position modulation*, i.e., to modulate the response delay Δ_{RESP}. For this purpose, we introduce $\Delta'_{RESP} = \Delta_{RESP} + \delta_i$, where δ_i sets an additional individual delay for each responder i. This reduces the probability of overlapping responses, as the latter are spread over a wider range of the CIR. In total, the CIR estimate provided by the DW1000 has a length of 1016 samples (for $PRF = 64MHz$) with a sampling period of $T_s = 1.0016ns$. Thus, the maximum offset is $\delta_{max} \approx 1017ns$, which relates to a maximum distance offset of $\delta_{max} \cdot c \approx 307m$. Knowing the maximum communication range and an estimate of the delay spread allows to define the number of non-overlapping responses fitting in the CIR.

VIII. Combining Response Position Modulation with Pulse Shaping for a higher scalability

The response position modulation proposed in Sect. VII poses a stringent limitation on the maximum number of users due to its dependency on the communication range r_{max}. Indeed, considering that the range of UWB transceivers can easily extend to $r_{max} > 75$m [7], [14], just up to $N_{RPM} = \frac{\delta_{max} \cdot c}{r_{max}} \approx 4$ responders are supported to ensure non-overlapping responses. Due to this limitation, we suggest to combine response position modulation with pulse shaping to increase the maximum amount of users that can make simultaneous use of concurrent ranging, and hence its scalability.

To this end, we use response position modulation to split the channel impulse response into N_{RPM} *slots* separated by $\delta = \frac{\delta_{max}}{r_{max}}$, hence reducing the number of responders which possibly interfere with each other. This allows to mitigate the impact of overlapping responses and multipath reflections. Within each slot, we use pulse shaping to identify the responder, as described in Sect. V. The number of pulse shapes N_{PS} defines the number of supported users per slot. Each responder is assigned to a slot as well as a pulse shape depending on its responder ID. The used slot is $n_{RPM} = ID \ \% \ N_{RPM}$, whilst the used pulse shape is $n_{PS} = \lfloor ID/N_{PS} \rfloor$. Depending on the slot number n_{RPM}, we set the additional delay $\delta_i = n_{RPM} \cdot \delta$.

In Figure 8 we divide the CIR into $N_{RPM} = 4$ slots using response position modulation. Within a slot, up to $N_{PS} = 3$ responders are active. In the example in Figure 8, three responders make use of slot 0 and slot 3, two responders make use of slot 1, and one responder makes use of slot 2. Therefore, the total number of responders is $N = 9$. The maximum number of responders that can make use of concurrent ranging in this scenario is $N_{max} = N_{RPM} \cdot N_{PS} = 12$. One can increase N_{max} (and hence the scalability of the system) by increasing the amount of pulse shapes N_{PS} and by increasing the number of slots N_{RPM}. As discussed in Sect. V the maximum number of pulse shapes is approximately 100. Assuming that the communication range r_{max} is limited manually by adapting the physical layer settings to 20m (which is sufficient for typical indoor applications [15]), the number of supported responders becomes more than 1500. Thus, using the proposed technique, the initiator requires just a single transmit and receive operation to estimate the distance to all neighbors simultaneously. Using classical two-way ranging schemes, instead, the initiator would need to send a packet to and to receive a packet from all the 1499 neighbors, respectively. This emphasizes the impact of the presented solutions, especially when having a high number of neighbors.

IX. Conclusions and Future Work

In this paper, we developed a practical solution for concurrent ranging enabling an efficient distance estimation to multiple users in parallel. We first proposed an algorithm to detect responses within the estimated channel impulse response automatically, and showed that this technique is highly-effective also in the presence of overlapping responses. Second, we encoded the responder ID in the CIR using different pulse shapes, making it possible to assign distance estimations to specific responders. Third, we introduced a technique called response position modulation to mitigate the impact of strong multipath components. Finally, we combined response position modulation and pulse shaping to increase the number of supported users and the scalability of the proposed concurrent ranging scheme.

In future work, we plan to use concurrent ranging to build an efficient cooperative or anchor-based localization system. Furthermore, in the work presented in this paper, we have neglected the impact of non-line-of-sight situations on the performance of concurrent ranging. We will hence investigate this impact thoroughly in the next months.

Acknowledgments

This work was performed within the LEAD-Project "Dependable Internet of Things in Adverse Environments", funded by Graz University of Technology (Graz, Austria).

References

[1] M. Z. Win and R. A. Scholtz, "Impulse radio: How it works," *IEEE Communications letters*, vol. 2, no. 2, pp. 36–38, 1998.

[2] D. Porcino and W. Hirt, "Ultra-wideband radio technology: potential and challenges ahead," *IEEE Communications Magazine*, vol. 41, no. 7, pp. 66–74, Jul. 2003.

[3] P. Catherwood and W. Scanlon, "Ultrawideband communications - an idea whose time has still yet to come?" *IEEE Antennas and Propagation Magazine*, vol. 57, no. 2, pp. 38–43, May 2015.

[4] IEEE, *Standard for Local and metropolitan area networks. Part 15.4. Amendment 1: Add Alternate PHYs*, Std. 802.15.4a, 2007.

[5] Decawave Ltd., *DW1000 Datasheet. Version 2.12*, 2016.

[6] J. Zhang, P. V. Orlik, Z. Sahinoglu, A. F. Molisch, and P. Kinney, "UWB Systems for Wireless Sensor Networks," *Proceedings of the IEEE*, vol. 97, no. 2, Feb. 2009.

[7] B. Großwindhager, M. Rath, C. A. Boano, and K. Römer, "Enabling Runtime Adaptation of Physical Layer Settings for Dependable UWB Communications," in *Proceedings of the* 19^{th} *IEEE International Symposium on a World of Wireless, Mobile and Multimedia Networks (WoWMoM)*, Jun. 2018.

[8] J. Kulmer, S. Hinteregger, B. Großwindhager, M. Rath, M. S. Bakr, E. Leitinger, and K. Witrisal, "Using Decawave UWB Transceivers for High-accuracy Multipath-assisted Indoor Positioning," in *Proc. of the Intern. Conf. on Commun. (ICC) Workshops*, 2017, pp. 1239–1245.

[9] B. Großwindhager, M. Rath, J. Kulmer, S. Hinteregger, M. S. Bakr, C. A. Boano, K. Witrisal, and K. Römer, "UWB-based Single-anchor Low-cost Indoor Localization System," in *Proceedings of the* 15^{th} *ACM International Conference on Embedded Networked Sensor Systems (SenSys), demo session*. ACM, Nov. 2017.

[10] P. Corbalan and G. P. Picco, "Concurrent Ranging in Ultra-wideband Radios: Experimental Evidence, Challenges, and Opportunities," in *Proceedings of the 15th International Conference on Embedded Wireless Systems and Networks (EWSN)*, Feb. 2018.

[11] Decawave Ltd., *DW1000 User Manual. Version 2.10*, 2016.

[12] IEEE Computer Society, *Standard for Low-Rate Wireless Networks*, Std. 802.15.4, 2015.

[13] C. Falsi, D. Dardari, L. Mucchi, and M. Z. Win, "Time of Arrival Estimation for UWB Localizers in Realistic Environments," *EURASIP Journal on Advances in Signal Processing*, no. 1, Jul. 2006.

[14] Decawave Ltd., *Application Note APS017. Maximising Range in DW1000 Based Systems. Version 1.0*, 2014.

[15] D. Dardari, P. Closas, and P. M. Djurić, "Indoor Tracking: Theory, Methods, and Technologies," *IEEE Transactions on Vehicular Technology*, vol. 64, no. 4, Apr. 2015.