Bidirectional TSCH for swarms of mobile robots in GNSS-denied environments

Edouard de Langlais
LAAS-CNRS, CNRS,
Université de Toulouse,
Toulouse, France
edouard.de-langlais@laas.fr

Lenny Barazzuoli
LAAS-CNRS, CNRS,
Université de Toulouse,
Toulouse, France
lenny.barazzuoli@laas.fr

Nicola Accettura
LAAS-CNRS, CNRS,
Université de Toulouse,
Toulouse, France
nicola.accettura@laas.fr

***Abstract*—The motion control of robots within swarms depends on Global Navigation Satellite Systems (GNSS), which could be unavailable in critical environments such as underground facilities, dense urban areas, or zones of conflict. An alternative way for robots to assess their own position in the swarm is offered by the Two-Way Ranging (TWR) methods over the IEEE 802.15.4 Ultra-Wideband (UWB) physical (PHY) layer, while the IEEE 802.15.4 Time-Slotted Channel Hopping (TSCH) provides a robust, time-deterministic medium access that the inter-robot commands require. Yet, TSCH supports neither double-sided radio ranging nor the bidirectional traffic that symmetric cooperative localization demands. To address all these requirements, we propose two modified timeslot structures, 3B-TSCH and 4B-TSCH, that fold a bidirectional acknowledged exchange and a double-sided ranging into a single cell. We present a preliminary physically grounded Monte-Carlo simulation-based evaluation of a Bidirectional TSCH link for mobile robots, where both communication and positioning rely solely on a new timeslot structure and TWR-enabled UWB. The results show that the proposed modification effectively improves the delivered information per unit of energy while adding ranging capability to mobile robots with the final goal of making robot swarms run safely in outdoor GNSS-denied environments.**

***Index Terms*—TSCH, mobile robots, UWB, Localization, GNSS-denied environments.**

I. Introduction

Autonomous robots, and more specifically their organization into cooperative swarms, have become an increasingly active research topic for both academia and industry [1]. They promise wider coverage, higher robustness and greater operational flexibility than a single high-end device, with applications ranging from entertainment drone light shows to environmental monitoring, search-and-rescue and military operations. Specifically, the coordinated motion of such swarms relies on the ability of each robot to estimate its own position, and this is usually done through a transceiver able to capture signals from Global Navigation Satellite Systems (GNSS). However, this technological dependency becomes a critical weakness in GNSS-denied environments, such as underground facilities, dense urban canyons or zones of conflict, where the satellite signal is attenuated, reflected, or deliberately jammed. Deprived of an absolute position reference, individual robots would accumulate drift in their position estimate, and the entire swarm would rapidly lose the spatial consistency required to navigate safely and maintain its inner formation. Several research investigations propose to mitigate this limitation by equipping each robot with onboard perception, e.g., cameras or Light Detection And Ranging (LiDAR) capabilities coupled with Simultaneous Localization and Mapping (SLAM) and Inertial Measurement Units (IMUs) [2]. Although effective, these solutions impose a significant burden in terms of computational load, energy consumption, and hardware cost, which is hardly compatible with the limited budget and endurance of large swarms of small drones.

An attractive alternative is to allow the robots to estimate their relative position directly from the radio links they already maintain to communicate, thereby reusing the communication infrastructure for localization instead of adding dedicated sensors. This idea has motivated cooperative localization schemes in which mobile nodes range against each other to recover the geometry of the swarm. Among the available radio technologies, Ultra-Wideband (UWB) stands out as a particularly well-suited candidate due to its accuracy in ranging at centimeter level, its license-free operation, its high data rate, and its energy efficiency [3] [4]. UWB is defined as a physical layer (PHY) in the IEEE 802.15.4 standard [5]. Beyond UWB, the standard also defines a radio ranging technique between pairs of devices, namely Two-Way Ranging (TWR), that can be implemented with either just one initiator node ranging its distance to a responder node (Single-Sided TWR, SS-TWR) or enabling both sides to range their own mutual distance (Double-Sided TWR, DS-TWR). Clearly, only the latter fits the cooperative localization scheme envisaged by this contribution. Building on these properties, several works have envisioned UWB as the backbone positioning system of mobile robot swarms [6]–[8], often fused with IMUs data and known anchor nodes to improve robustness. Unfortunately, UWB does not natively define how the medium should be shared among many devices: its impulse-radio modulation is poorly suited to carrier-sensing schemes [9], so that uncoordinated access quickly degenerates into ALOHA-like contention, with collisions and a sharp drop in throughput and reliability as the swarm grows. Therefore, scheduled medium access is an extremely wanted feature to fully exploit UWB in dense and mobile deployments, because of the inherent possibility to schedule collision-free communications.

For this purpose, IEEE 802.15.4 Time-Slotted Channel

979-8-3195-0167-7/26/$31.00 ©2026 IEEE

Hopping (TSCH) is a well-established low-power Medium Access Control (MAC) protocol tailored for collision-free communications over multi-hop sensor meshes. In detail, it controls medium access by letting devices communicate over a synchronized structure while following a communication schedule. In this way, the computation of collision-free schedules is delegated to higher layers [10]. Combined with appropriate routing and scheduling, TSCH thus yields a deterministic and highly predictable behavior for time-sensitive industrial applications [11], achieving high packet delivery ratios over multi-hop paths between sensor nodes and a single collection point. In this sense, scheduling collision-free UWB-based communications through TSCH would be a natural choice for swarms of mobile robots. It should be noted that TSCH implementations and scheduling algorithms are designed for static topologies and remain highly sensitive to the application scenario [12], [13]. Therefore, the implementation of a classic TSCH dedicated timeslot on top of the UWB physical layer would natively provide a transmitting node (TX) with the capability to range its distance from the associated receiving node (RX) within that timeslot, using the SS-TWR method. Instead, to allow these 2 neighbor nodes to mutually range their distance through DS-TWR, they should schedule an additional dedicated timeslot with the TX and RX roles inverted. Notably, TX and RX timeslots are independently scheduled on a node. While such a feature is extremely suitable for static TSCH networks (nodes schedule more TX timeslots than RX ones towards their parents), traffic flows into mobile swarms are less predictable, as the relative position of robots with respect to some collection points can vary over time. Thus, neighboring robots must interact through fully bidirectional data exchanges, i.e., they should install 2 complementary timeslots, one per direction. However, this first envisaged solution comes with 2 unwanted behaviors. First, if the minimum time interval between the 2 timeslots is sufficiently large, the 2 SS-TWR processes will produce 2 different distance estimates when the 2 neighbors move relative to each other. Such an impairment would be quickly reflected in the wrong control of the swarm at the higher layers. Second, scheduling the 2 timeslots adjacently would double the minimum end-to-end delay for a multihop delivery [10], [14], while the ACK of the first timeslot and the data frame of the second hold a redundancy that wastes both energy and throughput.

To cope with these potential issues, we propose a new TSCH timeslot structure able to accommodate a bidirectional and acknowledged data exchange between an initiator and a responder. Such a new dedicated timeslot is installed on 2 neighbors and allows: (i) one data delivery from the initiator to the responder; (ii) one data delivery from the responder to the initiator, piggybacked with the ACK to the previous exchange; (iii) the ACK of the latest delivery sent by the initiator to the responder. This timeslot natively supports the basic DS-TWR ranging method and bidirectional data delivery at the same time. In the remainder of this paper, we will refer to this 3-way Bidirectional TSCH structure as 3B-TSCH. For the sake of completeness, we have also considered the possibility to further extend the timeslot structure to support a fourth short data frame delivery from the responder to the initiator, in order to report the distance as measured by the responder and allow both devices to get a better distance estimate. We will refer to this 4-way Bidirectional TSCH case as 4B-TSCH.

We have evaluated the performance of the conceived solutions in terms of ranging error and energy consumption through simulations. Specifically, we have made a preliminary analysis that involves the communication between 2 devices in an outdoor scenario, and paves the way for future investigations of connectivity within a swarm. Both double-sided proposals hold a millimeter-level accuracy and a drift-free estimate under motion, while the single-sided counterpart stays bound to the tens-of-centimeters range, degrading to meters on a commodity crystal. A trade-off is also exposed between the two variants: while 3B-TSCH delivers the most information per unit of energy, 4B-TSCH spends one extra frame to grant both devices the drift-canceled estimate.

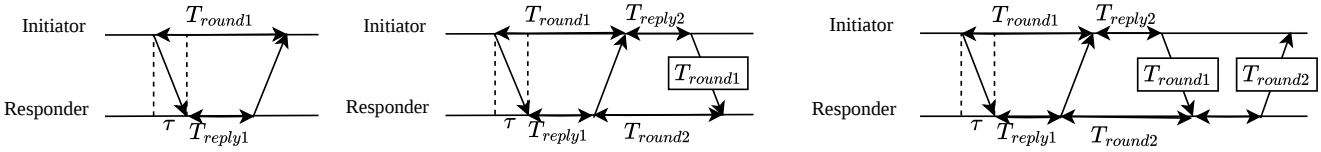
After having provided a sufficient technological background in Sec. II, we describe in great detail the 3 previously introduced Bidirectional TSCH implementations in Sec. III. Then, Sec. IV describes the methodological approach used to compare these solutions, while Sec. V presents and discusses the results of such a preliminary analysis. Finally, Sec. VI concludes the paper and envisages future work.

II. Technological background

The purpose of this section is to introduce the necessary technological details of the considered physical (Sec. II-A) and MAC (Sec. II-B) layers and help the reader figure out the scope of this contribution. Then, Sec. II-C introduces some related attempts to deal with localization through TSCH.

A. UWB physical layer

IEEE 802.15.4 Ultra-wideband is a communication technology defined in 2 PHY layers, but we focus on the High Rate Pulse (HRP-UWB) one, which features high data-rates and ranging accuracy, thus making it the best candidate for mobile robot swarms. A UWB frame can be transmitted over one among 16 channels defined over 2 bands (0 GHz to 3.9 GHz and 6.4 GHz to 9.9 GHz). Specifically, a UWB frame is composed of 4 fields. The Synchronization (SYNC) field holds the preamble used by the receiver to lock onto the message. Then, the Start of Frame Delimiter (SFD) code indicates the precise start of the Physical Header (PHR), used for timestamping the reception of the UWB frame and performing range estimation. After that, the PHR field carries the frame settings and error correction code. Finally, the PHY layer payload follows. In such a configured framework, the ranging capability is enabled by the modulation, with bursts of 2 ns encoded pulses being sent over a 499.2 MHz bandwidth or greater. Each symbol carried by a burst is emitted at a mean Pulse Repetition Frequency (PRF), which can vary between 62.4 MHz and 249.6 MHz. The generated pulses enable timestamping with a 15.65 ps resolution [5]. These



(a) Single-Sided TWR (b) 3-Way Double Sided TWR (c) 4-Way Double Sided TWR

Fig. 1: Two Way Ranging (TWR) variants in IEEE 802.15.4 [5]

timestamps can then be exploited by clock-drift-resilient radio-ranging techniques, such as DS-TWR [5].

As illustrated in Fig. 1, TWR relies on the exchange of 2 or more timestamped messages. In SS-TWR (Fig. 1a), the initiator sends a first frame. The responder starts transmitting a reply exactly after a T_{reply1} delay starting from the reception of the first frame. Such a delay is *a priori* known by both devices. The initiator measures the round-trip T_{round1} elapsed between the start of its own transmission and the start of the reception of the reply. So, the propagation delay is

$$\tau_{SS} = \frac{T_{round1} - T_{reply1}}{2}, \tag{1}$$

with the physical distance being estimated as $d = c\tau$, where c is the speed of light. It is worth noting that the value of T_{reply1} is the one measured by the clock of the initiator, while the original value was generated by the clock of the responder. So, the relative drift between the 2 clocks introduces a bias on the final estimate: SS-TWR is therefore inherently drift-limited. Instead, with 3-way DS-TWR (3DS-TWR) of Fig. 1b, the initiator sends a third frame to the responder after a T_{reply2} delay from the start of reception of the second frame. This frame also includes the measured roundtrip time T_{round1}. The responder can also measure the roundtrip T_{round2} and can compute an estimate of the propagation delay by using knowledge of all 4 intervals

$$\tau_{DS} = \frac{T_{round1}T_{round2} - T_{reply1}T_{reply2}}{T_{round1} + T_{round2} + T_{reply1} + T_{reply2}}. \tag{2}$$

This estimate is more accurate than (1), since the relative clock drift is canceled out [15]. With a fourth frame reporting T_{round2} from responder to initiator, the 4-way DS-TWR (4DS-TWR) of Fig. 1c permits both devices to accurately range each other according to (2).

B. TSCH medium-access layer

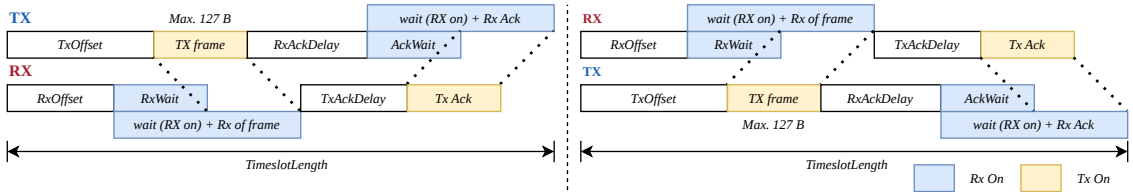
TSCH was initially defined within the Time Synchronized Mesh Protocol [14], and then standardized within the IEEE 802.15.4 specification [5]. It is based on the tight synchronization of all nodes in a mesh to a shared understanding of time. To achieve that, TSCH divides time into fixed equal-length timeslots, and a fixed set of timeslots forms a slotframe that repeats indefinitely. As depicted in Fig. 2a, each timeslot accommodates the transmission of a single frame of up to 127 bytes followed by a 43-byte ACK, both bounded to precisely timed transmission and reception windows. This tight timing enables the propagation of a single clock reference across the whole network, taken either from the frame transmitter (frame-based synchronization) or from the acknowledging receiver (ACK-based synchronization). In both cases, a device estimates the offset to its time source from the difference between an expected and an observed frame arrival time. The standard TSCH treats this difference as a pure clock offset, although it also contains the frame propagation delay. Knowledge of the latter becomes important if we want to exploit the radio ranging capabilities of UWB frames for the configured system.

In addition, TSCH runs over a set of logical channels, thus enabling concurrent non-interfering transmissions within the same timeslot boundaries. Specifically, when a TX node needs to transfer data to an RX neighbor, a "cell" should be scheduled on both nodes, meaning that they will both reserve the pair {timeslot, logical channel} within the slotframe. Remarkably, TSCH implements channel hopping by mapping the logical channel to a different physical one timeslot by timeslot, so that if the slotframe length is coprime with the number of channels, transmissions fairly span all physical channels and gain robustness against interference and multi-path fading. Although classic TSCH runs on top of the 16 available channels provided by the IEEE 802.15.4 O-QPSK 2.4 GHz DSSS physical layer [5], in the context of this contribution, TSCH is forced to run on top of the 3 selected channels introduced later in Sec. IV.

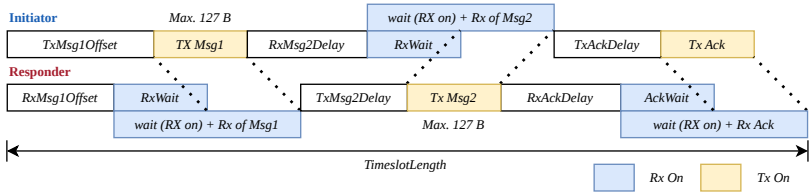
C. Previous work on the integration between TSCH and distance ranging

An attempt to combine TSCH with DS-TWR over UWB communications was already done in [16] by fitting the exchange of 4 messages into a TSCH timeslot. This attempt assumed 2 types of TSCH timeslots, both having the same duration: a classic TSCH data slot, with the same configuration as in subsection II-B; a "ranging slot," which includes the exchange of the 3 messages required by DS-TWR followed by a report message, for a total of 4 messages. As a matter of fact, from an architectural point of view an asymmetrical behavior is assumed: anchor nodes range mobile nodes. Instead, with the solutions proposed in the remainder of this contribution, any node can range other nodes and can be ranged as well.

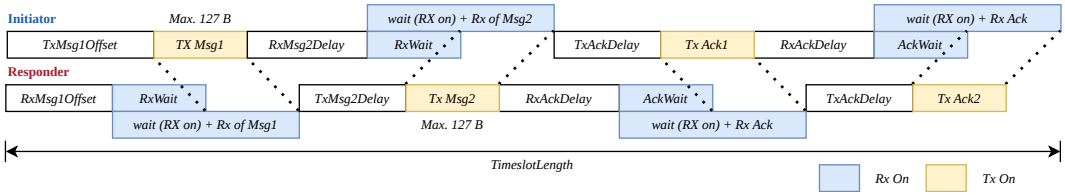
Other works keep communication and ranging in separate slots. In [17], TSCH schedules the distance measurements between mobile robots over a narrowband radio, so the UWB timestamping accuracy cannot be reached. More generally, TSCH scheduling over UWB remains highly sensitive to the deployment scenario [12], [13]. In contrast, our proposal



(a) Standard TSCH timeslot, per IEEE 802.15.4 2024 [5]



(b) 3B-TSCH timeslot



(c) 4B-TSCH timeslot

Fig. 2: Timeslot structures of standard TSCH and of the two proposed Bidirectional TSCH strategies.

merges bidirectional data exchanges and double-sided ranging into a single timeslot, turning every scheduled communication opportunity into a drift-canceled ranging opportunity.

III. Bidirectional TSCH

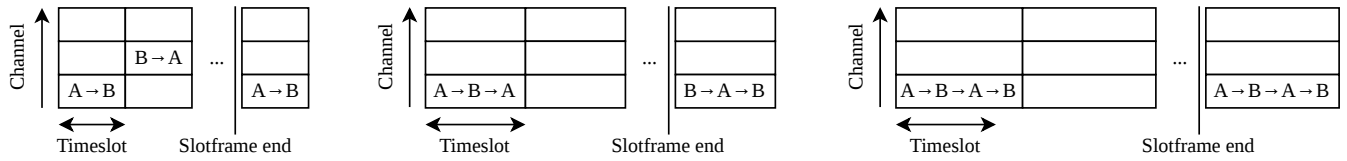
In this section, we introduce our approach to Bidirectional TSCH by defining two new slot structures, allowing native support for better ranging techniques and bidirectional data exchange. Yet, before introducing the dedicated slots, it is worth detailing how a bidirectional and ranging-capable exchange can be obtained with an *unmodified* TSCH timeslot, whose structure is recalled by Fig. 2a and whose schedule is shown in Fig. 3a. A standard TSCH timeslot is inherently unidirectional: carrying a single data frame from a TX node to an RX one, followed by an ACK. Therefore, letting the considered nodes send data *and* range each other at the same time requires the installation of 2 complementary slots with inverted roles, each allowing SS-TWR measurement. However, this solution would raise an issue from the perspective of end-to-end delivery: the airtime and the minimum multi-hop delay are inherently doubled, thus reducing the scope of time-critical applications. Moreover, the ACK of the first timeslot and the data frame of the second one flow in the same direction, while carrying redundant information, i.e., headers, trailers, etc. Such a configuration thus translates into a waste of energy and throughput. All these observations directly motivate the 2 timeslot configurations introduced hereafter. The 3B-TSCH configuration of Sec. III-A merges the 2 adjacent timeslots into a single one allowing the exchange of 3 MAC frames; instead, the 4B-TSCH configuration of Sec. III-B keeps 4 MAC frames within a single timeslot with appropriate changes. The timings of these new slot configurations are inspired by [9].

A. 3B-TSCH timeslot

The design goal of 3B-TSCH is to let a single timeslot allow a bidirectional acknowledged data exchange with 2 127 B messages and a 43 B ACK. These 3 messages are used to perform 3DS-TWR, as illustrated by the timeslot structure of Fig. 2b and its schedule in Fig. 3b. Thus, within one slot, the initiator transmits a first data frame (`Msg1`, up to 127 B) to the responder; after a fixed and known reply delay, the responder sends its own data frame (`Msg2`, up to 127 B), piggybacking the ACK of the previous data frame delivery; finally, the initiator sends back an ACK (`Ack`, 43 B) that additionally carries the roundtrip time measured by the initiator, i.e., T_{round1} of Fig. 1b. These 3 timestamped messages provide exactly the 4 intervals required by the DS-TWR equation (2); yet, at the end of the slot, only the responder holds all 4 terms and computes a drift-canceled range, whereas the initiator can only produce the single-sided estimate of (1). Since both devices exchange the same amount of data, their roles inside the timeslot are not constrained by a specific direction of the traffic flows. Thus, inverting the roles of initiator and responder every slotframe would allow each device to compute an accurate distance estimate from its neighbor according to (2) once every 2 slotframes without compromising the schedule, as illustrated in Fig. 3b.

B. 4B-TSCH timeslot

The 4-message strategy removes the asymmetry within the timeslot by adding a second ACK from the responder to the



(a) Standard TSCH (b) 3B-TSCH (c) 4B-TSCH

Fig. 3: TSCH Implementations Schedules

initiator, turning the slot into the structure of Fig. 2c, with its schedule in Fig. 3c: both devices send one $127\,\mathrm{B}$ data frame and one $43\,\mathrm{B}$ ACK. The last ACK reports the roundtrip time measured by the responder, i.e., T_{round2} of Fig. 1c, thus allowing both sides to get a drift-canceled estimate of their mutual distance through (2) at the same time. 4B-TSCH is the natural choice when simultaneous accurate ranging is needed at the cost of one extra frame, a longer timeslot and thus a longer slotframe for the same payload as the 3-way variant. On the scheduling side, the role inversion between initiator and responder (as already used for the 3B-TSCH timeslot) is not strictly needed, since both devices already end each timeslot with a drift-canceled range estimation.

C. Impact on synchronization

Since both 3B-TSCH and 4B-TSCH add messages to the standard slot, their synchronization behavior changes accordingly, although both native TSCH mechanisms frame-based and ACK-based synchronization remain available. We generalize the synchronization relation as

$$\epsilon = T_{round} - T_{reply} - 2 \times \tau, \tag{3}$$

where ϵ is the clock offset of a device relative to its time source, T_{round} is the measured round trip time between the data frame transmission and the reply reception, τ is the measured propagation delay, and T_{reply} the known reply delay (either `TxMsg2Delay` or `TxAckDelay` depending on which frame serves as the time reference, see Fig. 2).

Since τ is now estimated through the drift-canceled DS-TWR rather than assumed, the residual offset ϵ isolates the genuine clock drift, letting each device correct its clock against its time source node in any timeslot and regardless of its role (initiator or responder). This holds only as long as τ itself is trustworthy: since the ranged terms are timestamped at different instants within the slot, relative motion during the exchange may affect τ. We will quantify this effect in Sec. V-A.

IV. Methodological Approach

To estimate the performance of Bidirectional TSCH, a Monte-Carlo simulation tool has been developed, with each iteration focusing on two-slotframe exchanges. This allows us to isolate the fundamental issues that the bidirectional slots are designed to address: how accurately a pair of devices can range each other under a realistic MAC schedule, and how that accuracy degrades with clock drift and relative motion over successive ranging opportunities. The involved UWB transceiver is the Qorvo DW3110 [18], compatible with the newest versions of the standard [5]. Moreover, Qorvo is widely known for its UWB chips among the research community [4], [8], [19], and is regarded as the default transceiver for most applications. To this end, we implemented the 3 Bidirectional TSCH strategies presented so far: standard TSCH performing SS-TWR as the baseline, the 3B-TSCH with role inversion each slotframe, and the 4B-TSCH. This work uses the $499.2\,\mathrm{MHz}$ channels 5, 6, and 9 belonging to the HRP-UWB high-band. Each strategy is evaluated on a pair of randomly positioned devices, using clocks and transceivers modeled on datasheet-grounded values, to simulate the radio link between them and the on-board clock timestamping, as they would occur in a real implementation. Also, a proper energy consumption model has been developed based on the transceiver datasheet [18] in order to compare the cost of each approach. These span 50×5000 independent Monte-Carlo runs over 50 seeds that re-draw the two devices with random characteristics.

A. Device and geometry model

The chosen simulation layout involves 2 devices, namely device A, positioned at $(0, 0, 0)$, and device B placed at a random distance $d \in \mathcal{U}(0,\ 40\ \ m)$, random azimuth $\theta \in \mathcal{U}(0,\ 2\pi)$, and random elevation $\phi \in \mathcal{U}(-\pi/2,\ \pi/2)$. For the sake of clarity, we point out that with $\mathcal{U}(\cdot)$ we refer to a uniform distribution. Both devices have a random speed value $v \in \mathcal{U}(5,\ 15\ \ m.s^{-1})$ multiplied by a normalized random direction $\hat{dir} = (v_x, v_y, v_z)/\|(v_x, v_y, v_z)\|$ with each component $v_n \in \mathcal{U}(-1,\ 1)$, thus resulting in a uniform linear movement. Each device timestamps its frames with a two-state Wiener clock (White Frequency Modulation, WFM, and Random Walk Frequency Modulation, RWFM) [20], [21] with a per-device deterministic frequency offset drawn randomly in $\mathcal{U}(-f_{off},\ f_{off})$ (expressed in ppm). Two clock profiles have been considered. One is based on the Rakon IT2200K TCXO, whose WFM allows the $1\,\mathrm{s}$ Allan deviation to match the datasheet $0.3\,\mathrm{ppb}$ and $f_{off} = 1\ ppm$. The other one is a commodity crystal with $f_{off} = 20\ ppm$ and a literature-approximated short-term stability [22]. It has to be noted that the limit is a common UWB maximum operating range over the simulated radio channels (see Sec. IV-B and [23]). Moreover, by multiplying the maximum propagation delay, i.e., $40\ m/c$, by the maximum speed of a drone, i.e., $15\,\mathrm{m\,s^{-1}}$,

we get a maximum displacement error of $2\mu m$, which can be neglected in the simulation.

Each run is performed with the clock either enabled or disabled. For each of these cases, the system has been evaluated either with or without simulated motion. In this way, we can isolate the effect of drift from that of motion on the ranging accuracy.

B. Radio exchanges model

The UWB physical layer is modeled so that its inherent characteristics and default behaviors appear in a truthful simulation: transmissions are time-based and structured on the 2024 standard [5], and the propagation channel follows the model of [23], in order to ensure that each simulated link is physically grounded. For each device to communicate over UWB, the modeled DW3110 tightly follows the standard and copes with real physical parameters, such as receiver sensitivity and clock drift. However, since the DW3110 datasheet specifies the receiver sensitivity for only two of the three channels, we linearly interpolated the value for channel 6 from the neighboring channels 5 and 9. Each frame, except for the payload size, shares the same configuration with a high PRF mode, $64\,\mathrm{sym}$ SYNC length with code 9, an $8\,\mathrm{sym}$ SFD with code 1, fast PHR and $6.81\,\mathrm{Mbit\,s^{-1}}$. No Scrambled Timestamp Sequence is used.

The RF chain of any UWB link is described in (4) following the channel model of [23], where the received power at a distance d derives from a deterministic log-distance path loss, with $n = 1.58$ the path-loss exponent, $PL_0 = 48.96\ dB$ the reference loss at the reference distance $d_0 = 1\ m$. In the simulation runs, the model used is CM9, which accounts for an open-outdoor, line-of-sight (LoS) environment.

$$P_{\mathrm{RX}}(d) = P_{\mathrm{TX}} - \left[PL_0 + 10\, n \log_{10}\left(\frac{d}{d_0}\right)\right], \qquad (4)$$

A per-frame log-normal shadowing term ($\sigma_S = 3.96\ dB$ for CM9) gates delivery against the receiver sensitivity, accounting for environmental fading; it is distinct from the receiver thermal (AWGN) floor, which is already embedded in the transceiver sensitivity. Since only two devices with controlled emission and reception timings are simulated, collisions are not modeled. Antennas are considered isotropic, and the transmission power spectral density is capped at the regulatory $-41.3\,\mathrm{dBm/MHz}$, even though UWB emissions stay well below this cap. With the implemented timings and chosen data, t_{SHR} holds for $73.269\,\mu\mathrm{s}$, so the maximum allowed propagation distance is $d = c \times (2Rx_{guard} + t_{SHR}) = 3330\ km$ with respect to [9]. On the ranging side, each recorded timestamp is quantized to $15.65\,\mathrm{ps}$ ($128 \times 499.2\ MHz$ [5]), so the simulation reflects the finite time-stamping resolution of UWB hardware. No other ranging error is injected, so the observed error depends only on 3 factors: the timestamp quantization, the two-state Wiener clock drift and the geometry change due to motion. The comparison of TSCH, 3B-TSCH and 4B-TSCH thus reflects only how each technique is impacted by each setting.

TABLE I: Slotframe length and slot durations per technique.

Technique	Slots/frame	Slot duration (ms)	Slotframe (s)
Std TSCH	574	2.569	1.4748
3B-TSCH	287	3.405	0.9772
4B-TSCH	287	4.241	1.2170

V. Results analysis

Throughout the runs, the primary metric to look at is the per-exchange ranging error, i.e., the difference between the estimate and the true distance at the exchange instant. The discussion tackles the analysis of the effect of drift and motion (Sec. V-A), and the robustness of oscillator quality (Sec. V-B). We then weigh the schemes on radio energy and information efficiency (Sec. V-C). The considered scenario adopts a slotframe length of 287 timeslots in the case of 3B-TSCH and 4B-TSCH (see Table I), and 3 channels are used. Such a configuration would allow a dedicated cell for each link in fully meshed networks including up to 42 devices. In the case of standard TSCH, the considered slotframe doubles the number of timeslots to account for two dedicated cells per link, just enabling fully bidirectional communication and ranging. It is worth pointing out that the resulting TSCH slotframe allowing fully bidirectional communications is longer than the 3B-TSCH and 4B-TSCH ones, thus producing a net decrease in ranging opportunities.

A. Effect of drift and motion

Fig. 4 reports the per-exchange ranging Root Mean Square (RMS) error under the high-stability TCXO profile. With the clock disabled and no simulated motion, all three techniques sit at the $15.65\,\mathrm{ps}$ quantization floor (≈ 1 mm RMS), confirming that no artificial ranging error is injected. Once drift is enabled, any single-sided fix (TX node in standard TSCH or initiator in 3B-TSCH) jumps to the scale of $\approx 10\,\mathrm{cm}$ RMS of the DW3110 datasheet [18], whereas the drift-canceled double-sided fix of (2) held by the responder in 3B-TSCH and by both devices in 4B-TSCH stays pinned at the millimeter floor: one extra message buys a two-orders-of-magnitude accuracy gain.

With drift canceled through (2), the residual error is just due to relative motion. Within a slot, the exchanged frames are timestamped a few milliseconds apart, so the DS-TWR combination of (2) leaks a residual proportional to this time multiplied by the relative velocity. In this sense, the error is bigger for the initiator in the 4B-TSCH timeslot ($1.57\,\mathrm{cm}$) than for the 3B-TSCH responder ($\approx 0.82\,\mathrm{cm}$) and 4B-TSCH responder ($\approx 0.78\,\mathrm{cm}$). Indeed, the 4B-TSCH initiator will compute its range estimation according to (2) only upon reception of the 4-message, i.e., some time after the actual measurement, thus revealing the effect of motion.

B. Clock quality impact on the ranging accuracy

Fig. 5 shows the worst-case condition (both drift and motion on) as a per-estimator error distribution for the two clock profiles. Single-sided error stays at the $10\,\mathrm{cm}$ scale for the

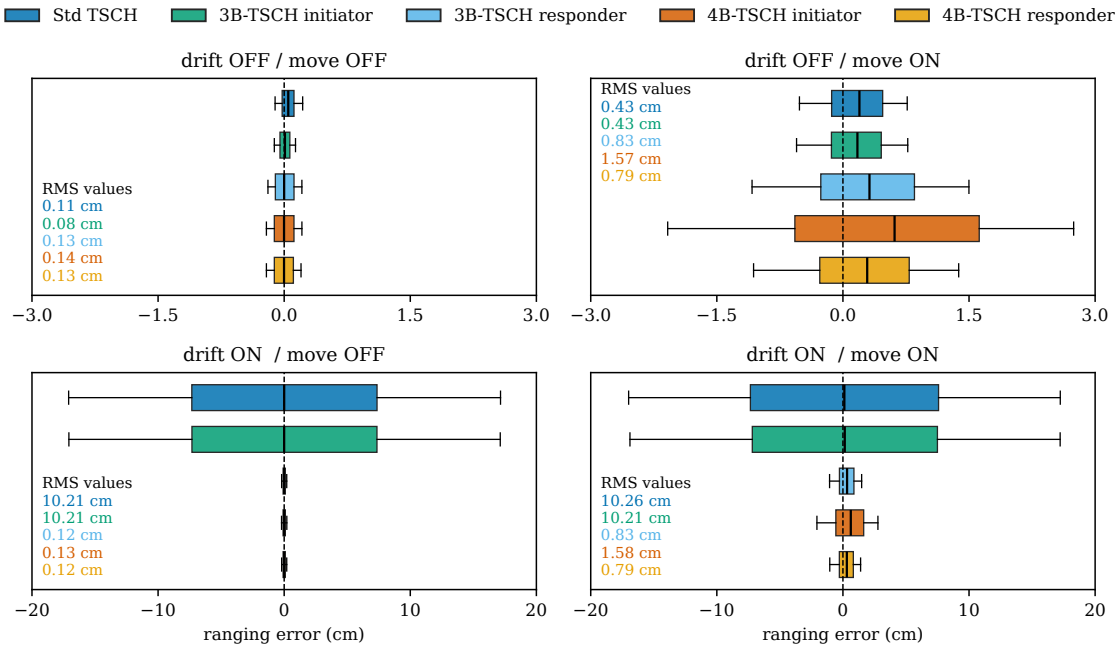


Fig. 4: Per-TSCH variant and conditions ranging error distributions under the TCXO.

TCXO, but degrades proportionally into the meter range for the literature XO, whose ± 20 ppm offset is 20 times larger. In contrast, the DS-TWR estimators demonstrate significant advantages for Bidirectional TSCH, maintaining centimeter-level accuracy on a sub-dollar crystal. Both devices reach it only in 4B-TSCH, however, since the 3B-TSCH initiator is left with a single-sided fix.

C. Energy cost and information efficiency

Ranging is only half of the equation: the extra frames of the Bidirectional TSCHs are not free in energy. To account for it, we model the active-radio cost of each exchange based on the DW3110 model from the datasheet [18]: the transceiver stays PLL-locked across a single active span from the first to the last frame, paying one wake-up, while the slot lead-in and tail are spent in deep sleep. Fig. 6a reports the resulting per-technique energy, stacked by radio state and split between the two roles of the exchange. Since more time is spent for listening and receiving, the cost of maintaining the RX state active is higher than actually transmitting. In standard TSCH and 4B-TSCH, the part related to the RX active state is higher than 3B-TSCH because of the 4 exchanges, while 3B-TSCH uses 3 data exchanges. Also, the 4B-TSCH suffers from a longer Idle PLL mode, where the transceiver is waiting for the next exchange but needs to keep its clock running, lacking time to get back to deep sleep.

Still, the energy consumed over a single timeslot does not show the full picture. In order to highlight which technique offers the highest information efficiency, we need to compute the information delivered over the amount of energy spent during a slotframe. This can be defined by

$$\eta_{slotframe}(tech) = \frac{254\ B}{E(tech) \times T_{slotframe}(tech)} \quad (5)$$

where $tech$ is the Bidirectional TSCH technique used, $\eta_{slotframe}$ is the efficiency per slotframe for the considered $tech$, 254 B is the amount of data exchanged in a single slot (or 2 slots for standard TSCH), E is the consumed energy according to $tech$ as already shown in Fig. 6a, and $T_{slotframe}$ is the slotframe size expressed in Table I. The results per technique are shown in Fig. 6b: while 3B-TSCH remains the most information efficient technique, 4B-TSCH also shows better performance than standard TSCH.

VI. Conclusion

We introduced 2 TSCH strategies that fold a bidirectional data exchange and a double-sided ranging into a single cell, so that every scheduled communication is also a drift-canceled ranging opportunity. Both were compared against standard TSCH through a physically grounded Monte-Carlo simulation of the UWB link, the device clocks and the radio energy. The study shows that drift cancellation is the decisive effect: single-sided ranging is fundamentally drift-limited, whereas the double-sided fix stays at the millimeter floor regardless of oscillator quality. 3B-TSCH delivers the most information per unit of energy, while 4B-TSCH better handles very high robot mobility. Once drift is canceled, the only remaining error comes from relative motion, which is mainly governed by the slot duration. More importantly, through this work we have shown that a proper modification of the TSCH timeslot can ensure the correct handling of bidirectional traffic in swarms of mobile robots.

This preliminary two-device analysis opens four research directions. First, the study will be scaled to full swarms, where dedicated scheduling and routing strategies must be designed to properly exploit spatial channel reuse. Then, future works will evaluate the impact of collisions and mobility on

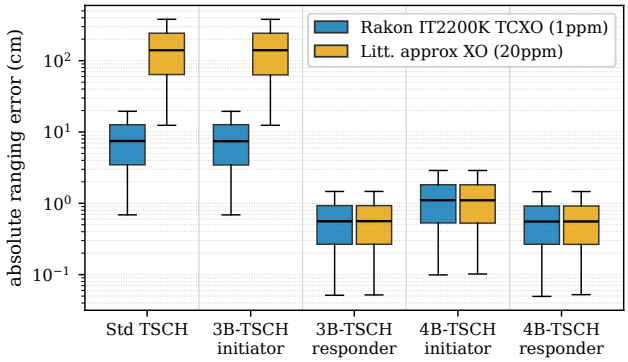


Fig. 5: Per-estimator ranging error for the two clock profiles, drift and motion enabled.

ranging latency and stability. Later, enhancing the estimation through Inertial Measurement Units (IMUs) or prediction-based filters like the Kalman or Particle ones is expected to further mitigate this residual motion error. Finally, the evaluation of full moving swarms will also be performed on a real DW3000/DW3110 testbed.

Acknowledgments

This work was co-funded by the French Defense Innovation Agency (AID) of the French Government Defense procurement and technology agency (DGA), and by the French National Research Agency (ANR) under the LAPSUS project (ref. ANR-23-CE25-0003).

References

[1] S. Javed *et al.*, "State-of-the-art and future research challenges in UAV swarms," *IEEE Internet of Things Journal*, vol. 11, no. 11, pp. 19 023–19 045, 2024.

[2] H. U. Unlu *et al.*, "Control barrier functions and LiDAR-inertial odometry for safe drone navigation in GNSS-denied environments," in *Motion Planning for Dynamic Agents*, 2023.

[3] Decawave Ltd, *DW1000 User Manual*, 2017.

[4] D. Coppens *et al.*, "An overview of UWB standards and organizations (IEEE 802.15.4, FiRa, apple): Interoperability aspects and future research directions," *IEEE Access*, vol. 10, pp. 70 219–70 241, 2022.

[5] *IEEE Standard for Low-Rate Wireless Networks*, IEEE Std. 802.15.4-2024 (Revision of IEEE Std 802.15.4-2020), 2024.

[6] Y. Zhou, Y. Wang, and Z. Zhong, "Study of mobile robot positioning based on ultra-wideband system," in *2022 IEEE 5th International Conference on Electronic Information and Communication Technology (ICEICT)*, 2022, pp. 595–598.

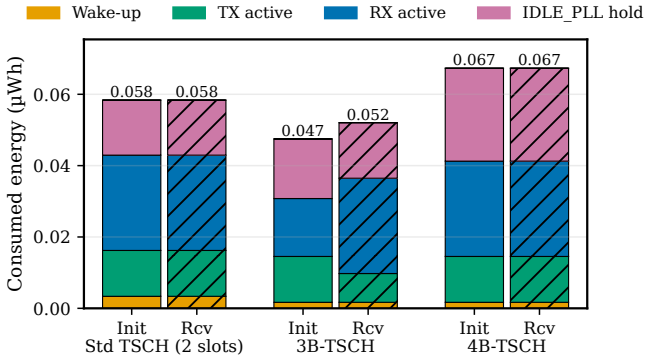
[7] K. Guo, X. Li, and L. Xie, "Ultra-wideband and odometry-based cooperative relative localization with application to multi-UAV formation control," *IEEE Transactions on Cybernetics*, vol. 50, no. 6, pp. 2590–2603, 2020.

[8] F. Shan *et al.*, "Ultra-wideband swarm ranging protocol for dynamic and dense networks," *IEEE/ACM Transactions on Networking*, vol. 30, no. 6, pp. 2834–2848, 2022.

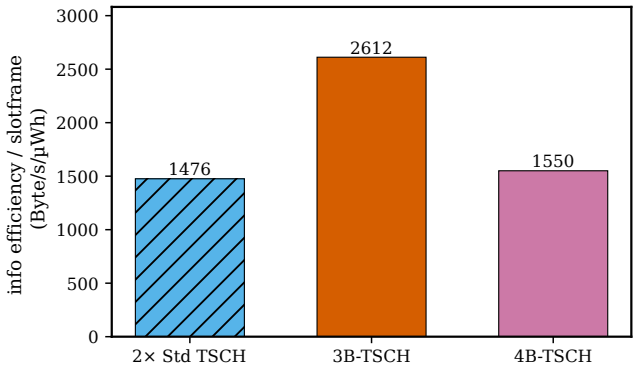
[9] M. Charlier, B. Quoitin, and D. Hauweele, "Challenges in using time slotted channel hopping with ultra wideband communications," in *Proceedings of the International Conference on Internet of Things Design and Implementation (IoTDI)*, 2019, pp. 82–93.

[10] N. Accettura *et al.*, "Decentralized traffic aware scheduling in 6tisch networks: Design and experimental evaluation," *IEEE Internet of Things Journal*, vol. 2, no. 6, pp. 455–470, 2015.

[11] M. R. Palattella *et al.*, "Standardized protocol stack for the internet of (important) things," *IEEE Communications Surveys & Tutorials*, vol. 15, no. 3, pp. 1389–1406, 2013.



(a) Energy consumption per method



(b) Slotframe-weighted delivered info efficiency

Fig. 6: Energy and information-efficiency comparison of the three techniques.

[12] C. Orfanidis *et al.*, "TSCH evaluation under heterogeneous mobile scenarios," *IoT*, vol. 2, no. 4, pp. 656–668, 2021.

[13] C. Orfanidis, P. Pop, and X. Fafoutis, "Active connectivity fundamentals for TSCH networks of mobile robots," in *2022 18th International Conference on Distributed Computing in Sensor Systems (DCOSS)*, 2022, pp. 191–198.

[14] K. S. J. Pister and L. Doherty, "TSMP: Time synchronized mesh protocol," in *Parallel and Distributed Computing and Systems*, 2008.

[15] D. Neirynck, E. Luk, and M. McLaughlin, "An alternative double-sided two-way ranging method," in *2016 13th Workshop on Positioning, Navigation and Communications (WPNC)*, 2016, pp. 1–4.

[16] M. Charlier, R.-A. Koutsiamanis, and B. Quoitin, "Scheduling UWB ranging and backbone communications in a pure wireless indoor positioning system," *IoT*, vol. 3, no. 1, pp. 219–258, 2022.

[17] G. Morano *et al.*, "Distance- and angle-based hybrid localization integrated in the IEEE 802.15.4 TSCH communication protocol," *Sensors*, vol. 24, no. 12, p. 3925, 2024.

[18] Decawave Ltd, *DW3000 Datasheet*, 2020.

[19] M. Wang *et al.*, "A 600-m UWB ranging dataset with radio frequency features and deep learning applications," *IEEE Sensors Journal*, vol. 25, no. 21, pp. 39 901–39 914, 2025.

[20] C. Zucca and P. Tavella, "The clock model and its relationship with the allan and related variances," *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control*, vol. 52, no. 2, pp. 289–296, 2005.

[21] L. Galleani, "A tutorial on the two-state model of the atomic clock noise," *Metrologia*, vol. 45, no. 6, p. S175, 2008.

[22] J. R. Vig, "Quartz crystal resonators and oscillators for frequency control and timing applications. a tutorial," U.S. Army Communications-Electronics Command, Tech. Rep., 1994.

[23] A. F. Molisch *et al.*, "IEEE 802.15.4a channel model - final report," IEEE 802.15 Working Group, Tech Rep. Document IEEE 802.1504-0062-02-004a, 2005.