Physical Limitations to the Accuracy of Range-based Localization and Sensing

Andrea Bedin*, Joerg Widmer*, Pablo Fernández Pérez,*† and Claudio Fiandrino*
*IMDEA Networks Institute, Spain, †Universidad Carlos III de Madrid, Spain
Email: {name.surname}@networks.imdea.org

***Abstract*—As we approach higher and higher localization and sensing accuracies with 6G technologies, the radio capabilities and signal processing will cease to be the main accuracy limitations, and other physical phenomena start to play an increasingly critical role. In this paper, we investigate how radio-based physical measurements, including radar, wireless localization and Integrated Sensing and Communication (ISAC), can be affected by environmental conditions, and at what level of accuracy and dimensional and temporal scales these effects become significant. Specifically, we produce and publish a dataset where the long-term stability of radio-based range measurements is empirically assessed, showing peak-to-peak variations of 0.3 mm over a 1 m range (which scales linearly, e.g. 3 mm over 10 m or 3 cm over 100 m) over 6 months. Further, we provide and substantiate a mechanistic explanation of these variations by showing how the combination of changes in refractive index of the air and thermal expansion of the room matches the measured data. We further demonstrate how such effects can be compensated for by measuring the atmospheric conditions and exploiting a combination of physical modeling and machine learning.**

***Index Terms*—Localization, Sensing, ISAC, 6G**

I. Introduction

Wireless sensing and localization are expected to be fundamental components of Sixth Generation (6G) networks [1], [2]. As technology advances, with larger antenna arrays, wider bandwidths, and increasingly sophisticated algorithms, positioning accuracy continues to improve, and has recently achieved millimeter-level accuracy using carrier phase information [3]. However, we will show that beyond a relative measurement accuracy of a few hundred of Parts per Million (ppm), the radio capabilities and signal processing techniques are not the only limitation, but significant contributions to the error can stem from other physical phenomena. This level of accuracy corresponds to a few millimeters over a 10 m range (in our specific conditions, around 3 mm), which could be of concern for precision manufacturing applications such as 3D printing with drones [4], or a few centimeters over a 100 m range, which might impact even less demanding tasks like the collection of small objects in a large warehouse.

In this paper, we provide the ISAC and radar communities with important insights on the magnitude and nature of external physical impairments on radio-based measurement, allowing other researchers to identify these effects, determine whether they are relevant in their work, and effectively address them.

This work received funding from the European Union, grant agreement 101201468 (MiRACLE) and 101129618 (UNITE), and by TUCAN6-CM (TEC-2024/COM-460), funded by CM (ORDEN 5696/2024). C. Fiandrino is a Ramón y Cajal awardee (RYC2022-036375-I), funded by MCIU/AEI/10.13039/501100011033 and ESF+.

To this end, we consider three key research questions:

How much do external physical effects impact radio-based range measurements? To *empirically* answer this question, in Sec. II we measure the stability of a radar range measurement over a long time and a variety of conditions, which in the conditions of this experiment are observed to be up to 3 mm over a 10 m range in a measurement period of 6 months. Albeit small, this error becomes significant for future carrier-phase based systems achieving millimeter accuracy [3]. Moreover, we will *publicly release* the firmware and hardware designs utilized in the measurements, as well as the raw dataset.

Which effects are the most impactful? We answer this question in Sec. III and IV. In particular, thanks to the well *controlled experimental conditions* (limited access room, clock frequency measurement and air temperature, pressure and humidity monitoring) we are able to model the building's thermal expansion and the changes in the air's refractive index, showing that they agree with the measured range variations. We note that while the individual effects are well-known and studied [5]–[7], the question of *which effects are practically relevant* in indoor ISAC and radar applications was not, to the best of our knowledge, previously addressed.

To what extent can we compensate for those effects? To address this question, in Sec. V we note that there are residual range variations that cannot be fully explained by the simple physical model. We determine that these variations are primarily due to the difference in temperature between the air and the concrete, and demonstrate that they can be explained with the measured environmental parameters by augmenting the prediction with Machine Learning (ML) models.

Furthermore, in Sec. VI we discuss the subtleties associated with thermal expansion, in Sec. VII we discuss the challenges associated with compensating these effects in realistic environments, and in Sec.VIII we conclude the paper.

II. Magnitude of the measurement variations

Let us begin by empirically answering the first question posed in Sec. I, i.e., *"How much do external physical effects impact radio-based range measurements?"*. To do so, we used four Acconeer XM125 radar boards [8] based on the A121 pulsed coherent radar chip [9], operating in the 57-64 GHz band. These radars were mounted on metal supports clamped to the ceiling of a storage room with monitored access. The choice of using a controlled environment, rather than a more realistic one, allows us to control for external interference and thus to correctly and reliably understand the mechanisms behind the

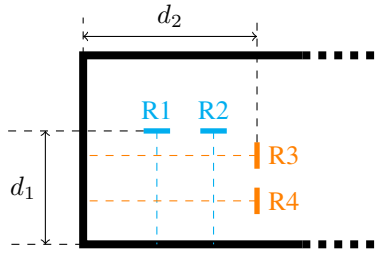

Fig. 1: Experimental setup

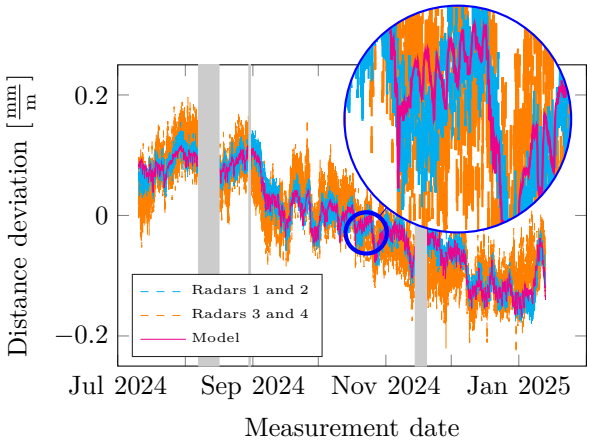

Fig. 2: Measured and modelled range deviation

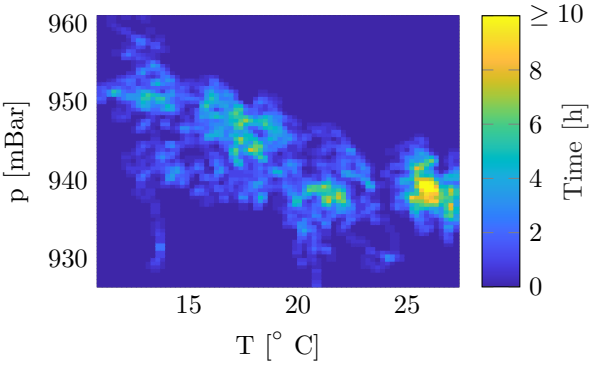

(a) Observed temperature and pressure conditions

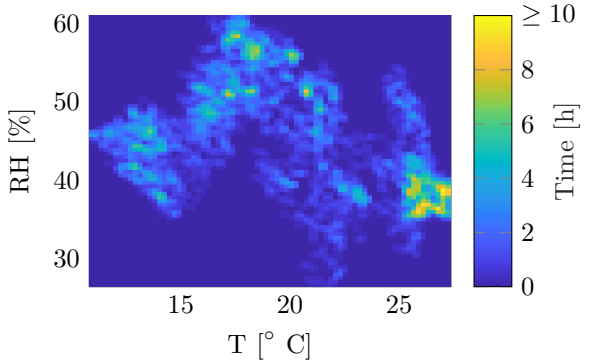

(b) Observed temperature and relative humidity conditions

Fig. 3: Atmospheric conditions observed over the whole experiment

observed changes. Further measurements to better determine how these effects behave in a more realistic environment are planned for future works. The radar's positions are illustrated in Fig. 1. Specifically, Radar 1 (R1) and Radar 2 (R2) were installed on the same support, facing a wall at a distance of approximately $d_1 \approx 5.7$ m. Meanwhile, Radar 3 (R3) and Radar 4 (R4) were aimed at a different wall, positioned about $d_2 \approx 2$ m away. To measure the range variations, we developed custom firmware for the XM125 boards, utilizing the sparse IQ service [10] to extract the phase of the reflection from the walls. The phase variations are then divided by 2π and multiplied by the wavelength to obtain the range variations. The hardware designs, firmware and dataset will be available at [11].

Fig. 2 shows the measured distance variations, with data from R1 and R2 shown in cyan and from R3 and R4 in orange (the magenta line will be discussed later). These values are averaged over 1 hour intervals and normalized by the measurement distances d_1 and d_2. Note that since $d_2 < d_1$, the noise appears more pronounced for R3 and R4. The experiment ran for over six months, with three interruptions due to logistical issues (power outages, PC crashes, etc...) indicated by gray bands in Fig. 2. Since the setup measures only distance variations rather than absolute distances, any break in data collection results in a discontinuity between the measurements before and after the interruption. Thus, to align these segments, the four data sections were shifted so that the average deviation matches the model in Sec. III, which demonstrated good agreement within each segment. The measurements show a peak-to-peak deviation of nearly $0.3\frac{\text{mm}}{\text{m}}$ (or 300 ppm) over 6 months, and variations up to $0.1\frac{\text{mm}}{\text{m}}$ within a month (e.g., September 2024). Note that the overlapping of the normalized traces of R1 and R2 with the one from R3 and R4 confirms that such variations scale linearly with the measured range. Thus, the variation over 6 months in a regular 10 m room can be in the few millimeters range, with centimeter level errors in a larger environment of many tens of meters. Imagining a future application like a swarm of 3D printing drones [4] manufacturing a large, complex structure that requires many days or months to complete, this might lead to a drift in drone positioning and thus a deformation in the structure.

III. Modelling the Range Variations

To address the second question posed in Sec. I, i.e., *"Which effects are the most impactful?"*, in this and the following section we assess the causes of the observed range variations. To this end, temperature, pressure, and Relative Humidity (RH) were measured every 43 ms using a BME280 sensor [12] configured according to the datasheet's recommendations for indoor navigation applications. Fig. 3 shows the set of conditions experienced during the experiment (from early July 2024 to mid January 2025). In particular, Fig. 3a shows the number of hours spent in each combination of temperature and pressure, whereas Fig. 3b shows the same for the combinations of temperature and RH. Here we can observe that a large portion of the conditions in the range of 10° to 28° of temperature, 920 hPa to 960 hPa of pressure and 20% to 60% relative humidity has been observed for a sufficient amount of time. Further, to measure the clock stability each radar was

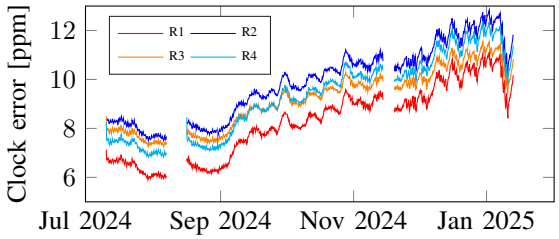

Fig. 4: Measured variations of the clock frequency

mounted on a custom PCB connecting its 24 MHz crystal oscillator to a TL3116CD comparator. This converted the low-amplitude oscillator's signal into a corresponding 24 MHz 3.3 V square wave without causing excessive load to the crystal. The comparator's output was used as external clock for Timer 1 of an STM32F411 Micro Contoller Unit (MCU) [13]. A 1 Pulse per Second (PPS) signal from a u-blox NEO-6M GNSS module [14], which provides a 1 Hz signal with an absolute error of < 60 ns (resulting in an absolute accuracy of 0.06 ppm), was used to trigger an interrupt on the MCU. At each interrupt, the Timer 1 counter value was read, reset, and logged, providing a precise measurement of the oscillator's frequency in Hz. The 0.06 ppm accuracy is more than sufficient to detect any significant change as the observed range variations in Fig. III are up to 300 ppm. The measured clock variations can be observed in Fig. 4, and are in the order of 5 ppm.

With this knowledge, we can compute the impact of the following physical phenomena:

A. Air refractive index changes

Wireless localization systems can not directly measure distance; instead, they measure time and convert it into range using the speed of light as a conversion factor. Typically, the speed of light in a vacuum c is assumed, since air has a low density and, therefore, a refractive index close to 1. While this approximation is generally accurate, the actual speed of light in air is approximately $\frac{c}{n}$, with $n \approx 1.0003$. This is often described by providing the difference with vacuum in ppm $N = \frac{n-1}{10^{-6}}$. Using this notation, we have that $N \approx 300$ on average, with variations up to roughly 100 ppm [6]. More in detail, N obeys (1), where T is the temperature in Kelvin, p_d the partial dry pressure and p_w the partial water pressure [6].

$$N = 77.6890\frac{p_d}{T} + 71.2952\frac{p_w}{T} + 375463\frac{p_w}{T^2}. \quad (1)$$

Atmospheric conditions are typically expressed in terms of total pressure[1] p in millibar and the relative humidity RH expressed in percentage of the vapor pressure. These values can be converted to p_d and p_w according to [15], which states that $p_d = p - p_w$, $p_w = \frac{1}{100} RH \cdot 6.1094 e^{\frac{17.625T'}{T'+243.04}}$, with T' the temperature in degrees Celsius. While changes in the order of 100 ppm might seem negligible, they can introduce significant errors at the high precision levels expected in 6G systems. For instance, consider a rainy summer day with $RH = 70\%$, $p = 990$ hPa, and $T = 310$ K ($\sim$ 37°C), resulting in $N = 417$. In contrast, on a dry winter day with $RH = 30\%$, $p = 1010$ hPa and $T = 280$ K ($\sim$ 7°C), we obtain $N = 294$, leading to a total change of 123 ppm. Over a 10 m distance, this difference results in a path length variation of $123 \text{ ppm} \times 10^{-6} \times 10 \text{ m} = 1.23$ mm.

[1]Sometimes atmospheric pressure is reported in meteorological datasets normalized by altitude. When using public data in this formula, the actual pressure should be used rather than the normalized one.

B. Thermal expansion

In most indoor sensing settings, the radars are mounted inside a room, which is inevitably subject to thermal expansion as the temperature of the concrete constituting the walls, floor and ceiling varies. While this is arguably a real change in range, and therefore measuring it might seem (and in some contexts is) desirable, this can lead to inconsistencies in the case of localization systems, that are better described in Sec. VI. The magnitude of this effect is comparable to that of refractive index variations. For instance, concrete has a Coefficient of Thermal Expansion (CTE) of approximately $\kappa = 10\frac{\text{ppm}}{K}$. Given a sensor and target separated by 10 meters, and using the same temperature change as in the previous example ($T = 280$ K in winter, $T = 310$ K in summer), their relative position would shift by $30 \text{ K} \times 10 \frac{\text{ppm}}{K} \times 10^{-6} \times 10 \text{ m} = 3$ mm.

C. Clock variations

Another potential source of error is the precision of the timing reference used to determine propagation delay. Most commercial devices rely on generic crystal oscillators for accurate timing, as opposed to the high accuracy temperature compensated and oven controlled oscillators used in measurement equipment. These oscillators exhibit small but non-negligible frequency variations due to factors such as temperature changes and aging. For example, let us assume that a carrier-phase range measurement is performed over a distance of 10m, and the oscillator used to generate the carrier exhibits an offset of $\epsilon_f = 100$ ppm. The resulting error in the measured distance would be $10 \text{ m} \times \epsilon_f \times 10^{-6} = 1$ mm. While some high-precision oscillators have frequency stability in the range of 10 ppm [16], making their impact an order of magnitude smaller than the effects of refractive index variations or thermal expansion, others have significantly higher variability [17], making their contribution comparable to other factors. Therefore, if the selected oscillator is sufficiently accurate, its impact can be ignored, but if the oscillator uncertainty exceeds a few tens of ppm its effect becomes comparable to the other phenomena, and accounting for it is crucial for achieving millimeter-level accuracy.

IV. Model Results

To compute the range variations, we select a reference distance of $d_0 = 1$ m and a reference condition of $T_0 = 20$ °C, $p_0 = 950$ hPa and $H_0 = 50\%$, leading to a reference refractive index $N_0 \approx 302.5$ ppm. Based on this value, and on the effects discussed in Sec. III, the predicted range variation with respect to this reference can be written (with a small approximation for the refractive index part) as:

$$\Delta d = d_0 \left[\kappa(T - T_0) - (N - N_0)\right] \times 10^{-6}. \quad (2)$$

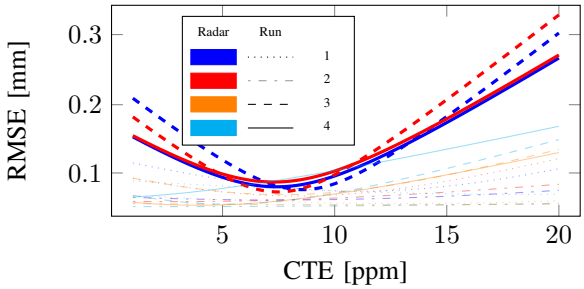

Fig. 5: Range RMSE vs CTE for different runs and radars

Clock frequency deviations are neglected here, as the observed clock variations are in the order of a few ppm (see Fig. 4), and are therefore negligible compared to the other effects. If the hardware is such that this effect is relevant, then $d_0 \Delta f$, with Δf the fractional deviation of the clock frequency, should be added to Δd. To determine the CTE κ of the room, we predicted the range variations observed by each radar using (2) for different values of the CTE, and compared them with the measured range variations. Fig. 5 shows the Root Mean Square Error (RMSE) of the predicted range variation for different values of the CTE for all runs and radars. Here we can see that R1 and R2 for runs 3 and 4 (thicker lines in Fig. 5) provide a much sharper minimum. This can be explained by two factors. First, R1 and R2 are positioned further from the wall compared to R3 and R4, which reduces noise influence when computing the RMSE. Second, runs 3 and 4 are considerably longer than runs 1 and 2 (where the experiment had to be interrupted due to logistical issues), and exhibit larger relative changes in environmental conditions. These larger variations produce more meaningful changes in range, which leads to a better CTE estimate. In view of this, we average the estimate of the CTE for those 4 cases to obtain a CTE of $7.5 \frac{\text{ppm}}{\text{K}}$.

The resulting values of (2) are shown as the magenta line in Fig. 2. Here, we can observe that the general trend matches quite well with the model, demonstrating that the considered effects are sufficient to explain most of the variations.

Fig. 6 presents the breakdown of the modeled range deviation into its two components: the refractive index variation in cyan and thermal expansion in orange[2]. From this figure, we note that thermal expansion generally contributes more to the overall deviation than the refractive index. Nevertheless, the contribution of the refractive index variations is still noticeable. This shows that both effects are indeed relevant, and cannot be neglected. Furthermore, we can see that in some instances, the two effects have opposite signs, partially canceling each other out and reducing the total deviation. However, in most cases, both components share the same sign, amplifying the overall effect and making the issue more pronounced.

V. Residual Variations

Despite, as shown in Sec. III, thermal expansion and refractive index changes explain *most* of the observed variations,

[2]Note that in some cases the two contributions have opposite signs, in that case the two areas overlap generating a darker cyan color.

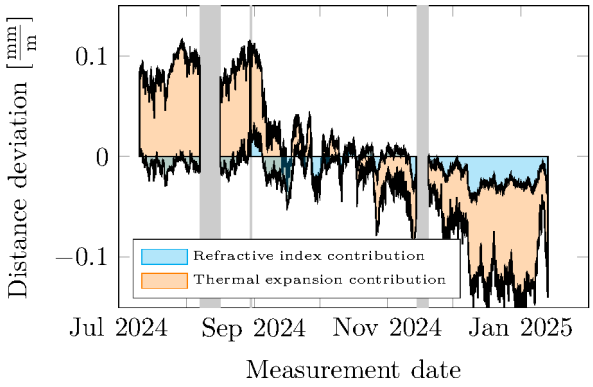

Fig. 6: Contribution of thermal expansion and refractive index to the range deviation

we note that some residual variations that do not seem to stem from random noise persist. In particular, the cyan trace in the top part of Fig. 7 shows the residual variations after removing the result of (2) from the data[3]. Here we can observe, e.g., a distinct negative bias around July 15$^{\text{th}}$ and daily oscillations in the order of $0.2 \frac{\text{mm}}{\text{m}}$ between the end of July and beginning of August (highlighted in red). This naturally brings up the third question posed in Sec. I: *"To what extent can we compensate for those effects?"*. In other words, can these residual variations be explained with the available data, and are therefore such variations a limitation of the model, or do they depend on other factors that are not measured in this experiment? To assess this, we first modeled the problem as a time-series prediction task and used a vanilla Neural Network (NN) (2 fully connected layers each with 400 hidden units, ReLU activation function). The output of the network is the predicted phase measured by the radar, while input is represented as a sequence $X \in \mathbb{R}^{L \times C}$, where $C = 3$ is the number of channels (temperature, pressure, and humidity) and $L = 144$ is the sequence length, corresponding to a 24 hour window at a 10 minute resolution in this evaluation. The data is split chronologically, with the first 60% of each run used for training and the remaining 40% as test set. Exploiting the universal function approximation capability of neural networks, we hope that the model learns a mapping from the physical model input features to the corresponding phase deviations.

The resulting prediction error is reported in Tab. I, where we can see that the vanilla NN underperforms even compared to the simple physical model. This is primarily due to two intrinsic limitations. First, in most time-series forecasting tasks, past values of the target variable serve as an endogenous input, providing a crucial inductive bias. In our case, however, the input cannot include past values of the ground truth itself, as it would not be known in a practical setting. Further, the data presents distribution drift, where the distribution of unseen data differs from that of the training set. Distribution drift is a well-known problem in ML, and particularly in time-series prediction. To exacerbate this problem, the absence of an

[3]Note that we show only a subset of the data for R1 and part of run 3 for readability, accuracy figures for the whole dataset can be found in Tab. I. Note also that the darker cyan color is the overlapping of the two traces.

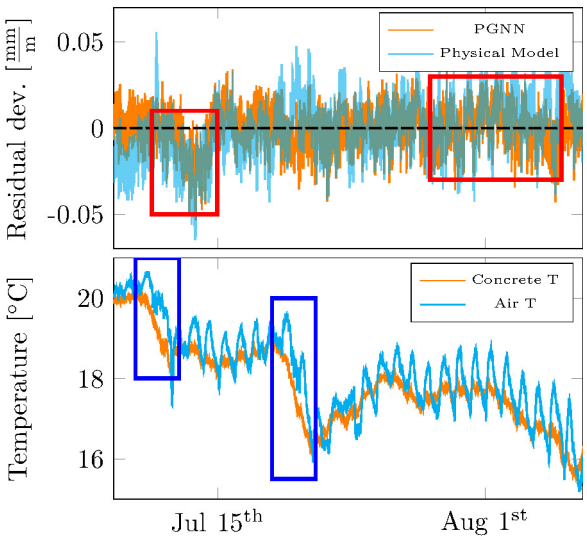

Fig. 7: Residual variations after compensating for thermal expansion and refractive index changes

endogenous variable prevents the use of techniques to mitigate distribution drift, such as reversible instance normalization [18], which normalizes input sequences by removing certain statistics (e.g., the mean) and later reintroduces them at the end of the forward pass.

To address the challenges described above, we explore an alternative solution. We note that a key limitation of the experiment is that we only measure air temperature, whereas it is the temperature of the concrete ceiling that directly drives thermal expansion. Based on this observation, we combine the domain knowledge of the physical model with the learning capacity of the NN, creating a hybrid approach in which the physical equations act as an inductive bias while the NN infers the actual wall temperature. This Physics Guided Neural Network (PGNN) is designed to predict the concrete temperature solely from historical values of the air temperature. The predicted temperature is then incorporated into the physical model as the effective temperature input for the thermal expansion. This physics-based approach enables the model to learn effective representations while avoiding the problems observed in the vanilla NN. The model architecture is identical to the vanilla NN, except that it outputs the temperature rather than the phase directly. With this insight, the ML model manages to outperform the simple physical model, as can be seen in Tab. I. Further, we note that the advantages of the PGNN over the Physical model are more pronounced in Runs 3 and 4, where larger variations in room temperature lead to stronger effects of thermal expansion and the overall influence of noise is diminished. Fig. 7 shows the residual deviations in the top subfigure for both the physical model and the PGNN, together with the concrete temperatures used in each case (equal to the air temperature for the physical model) in the bottom plot. It can be seen that the result obtained using the concrete temperature predicted by the PGNN provides a better

TABLE I: Distance deviation prediction RMSE [$\frac{\text{mm}}{\text{m}}$]. Reporting mean and standard deviation over five trials for NN and PGNN

Run	Model		
	Physical	NN	PGNN
1	0.0176	0.0309 ± 0.0005	0.0172 ± 0.0001
2	0.0188	0.0441 ± 0.0013	0.0181 ± 0.0001
3	0.0224	0.0492 ± 0.0020	0.0186 ± 0.0001
4	0.0319	0.0356 ± 0.0031	0.0221 ± 0.0001

prediction, which for example do not show the oscillations (highlighted in red) between the end of July and beginning of August. This is thanks to the fact that the concrete temperature remains more stable over time, consistently with the higher thermal mass of the concrete. In addition, when there is an abrupt change in temperature (highlighted in blue) we observe a lag between the temperature of the concrete and the temperature of the air. This is likely caused by the fact that the sun heats up the concrete first, and then the concrete transfers the thermal energy to the air slowly. Overall, these results show that a better prediction can be obtained by accounting for the difference in temperature between the concrete and the air. Thus, while the physical model correctly captures the mechanisms causing the observed variations, it needs to be applied with due care.

VI. Discussion on Thermal Expansion

As mentioned in Sec. III-B, one might argue that the thermal expansion constitutes a real change in range, and therefore is not an impairment.

In many situations, though, this can lead to inconsistencies. In this section, we want to elucidate this fact by laying out two, in our opinion reasonable, fundamental properties that a wireless localization or sensing system must satisfy and an example of their violation. For simplicity, we illustrate these concepts in one dimension. However, the same principles extend to 3D localization as well.

- P1: The measured range should change if and only if the target's position shifts relative to its local frame of reference. For example, if a robot's wheels do not turn, no movement should be detected in the range measurement. Moreover, any displacement of the target in its *local* frame of reference (e.g., how much a robot's wheels turn relative to the floor) should exactly correspond to the measured change in the range between the target and the sensor.
- P2: Property P1 must hold within a single, unambiguous, and well-defined global coordinate system. Specifically, if Property P1 is valid for one target in a given global coordinate system, it must hold for all other targets within the *same* global coordinate system.

Let us now consider a robot that drops a small object while recording its distance from a sensor (Fig. 8 (a)). Later, the robot attempts to retrieve the object, relying on the previously measured radio data. However, due to floor expansion caused by a temperature increase, the object's position with respect to the sensor changed, while it remained stationary with respect to the floor under it, violating Property P1. Let us imagine

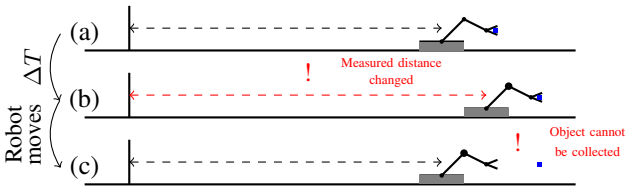

Fig. 8: Failure caused by thermal expansion

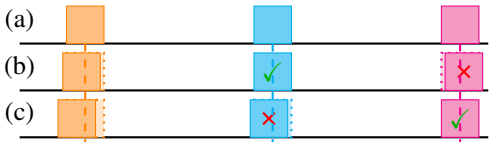

Fig. 9: Effect of thermal expansion on the reference frame

that the robot did not move through this process, but, e.g., it was rebooted and has no persistent memory, so it is unaware of this (Fig. 8 (b)). In this case, the robot would already be in the correct location, but based on the previously measured distance it would move to a different location, leading to a misalignment and a failed retrieval attempt (Fig. 8 (c)). Still, it could be argued that, since the robot does not move, and the distance changes, Property P1 could be still satisfied if we interpret the change as a movement of the sensor's location in the opposite direction. This would, however, lead to another contradiction. To illustrate this, consider a scenario with one sensor and two targets, where Property P1 is enforced by moving the sensor. Fig. 9 depicts the sensor in orange and the two targets in cyan and magenta. Case (a) represents the initial setup with no thermal expansion. Case (b) shows the system after a 2% thermal expansion, where Property P1 is enforced for the cyan target by shifting the reference frame (and consequently the sensor). Case (c) enforces Property P1 for the magenta target instead. Clearly, the global reference frames in (b) and (c) differ. Thus, if Property P1 is satisfied for one target, it no longer holds for the other, violating Property P2. For these reasons, in a setting where multiple targets are localized, thermal expansion can lead to inconsistent measurements.

VII. Challenges and Future Work

A natural question raised by these results is if and how we can mitigate these effects in a more realistic setting. A first approach would be to replicate the setup of this paper, measuring the environmental conditions and using the result of (2) to compensate. However, as highlighted in Sec.V, the concrete and air might not be in thermal equilibrium. While this leads to minimal variations in this controlled experiment, these discrepancies can be exacerbated by air circulation and proximity to Heating, Ventilation, and Air Conditioning (HVAC) systems. Further research is needed to verify that ML methods can still correctly predict the concrete temperature in such more complex settings. An alternative solution would be to use a fixed radio-based beacon and measure its displacement to estimate these effects. While this is in our opinion viable, it still comes with some caveats. First of all, in large rooms where these effects are more relevant the air conditions might not be uniform, thus, a single measurement might lead to an incorrect estimate. Secondly, buildings might be built unisotropically (e.g. with metal beams on a ceiling), and thus have a different CTE in different directions which will need to be measured independently.

VIII. Conclusions

Until recently, localization and sensing accuracy were primarily limited by radio and signal processing components. As the field advances, previously minor effects are becoming significant, requiring accurate modeling to ensure reliable and repeatable localization. This work shines a light on these phenomena with an extensive empirical evaluation, providing a quantitative understanding of the magnitude of these effects, which are shown to be relevant in high precision applications at long ranges. Further, we identified the main sources of error in thermal expansion and changes in air refractive index and showed that these effects can be modeled and compensated for in our setting. Finally, we highlight how differences in the temperature of the air and concrete can make this compensation a challenge, and demonstrate that an appropriate combination of ML and physical modeling can address this issue.

References

[1] Hexa-X, "Deliverable D3.1 - localisation and sensing use cases and gap analysis," https://hexa-x.eu/wp-content/uploads/2022/02/Hexa-X_D3.1_v1.4.pdf, 2021.
[2] S. E. Trevlakis *et al.*, "Localization as a key enabler of 6G wireless systems: A comprehensive survey and an outlook," *IEEE Open Journal of the Communications Society*, vol. 4, pp. 2733–2801, 2023.
[3] A. Bedin *et al.*, "Millimeter-scale absolute carrier phase-based localization in multi-band systems," 2025. [Online]. Available: https://arxiv.org/abs/2511.05204
[4] K. Zhang *et al.*, "Aerial additive manufacturing with multiple autonomous robots," *Nature*, vol. 609, no. 7928, pp. 709–717, Sep. 2022. [Online]. Available: https://doi.org/10.1038/s41586-022-04988-4
[5] L. Wang and B. Yazıcı, "Effects of fluctuation in refractive index of atmosphere on synthetic aperture radar images," in *2017 IEEE Radar Conference (RadarConf)*, 2017, pp. 0169–0174.
[6] J. RÜEGER, "Refractive index formulae for radio waves," *Proc. FIG XXII International Congress, Washington, D. C.*, 01 2002.
[7] M. Lazecky *et al.*, "Bridge displacements monitoring using space-borne x-band sar interferometry," *IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing*, vol. 10, no. 1, pp. 205–210, 2017.
[8] *Entry+ Module XM125 Datasheet v1.0*, ACCONEER, 03 2023.
[9] *A121 – Pulsed Coherent Radar (PCR)*, ACCONEER, 03 2023.
[10] *A121 Sparse IQ Service User Guide*, ACCONEER, 4 2024.
[11] https://softwarewng.networks.imdea.org/carrier-phase-stability-measurement/, 2025.
[12] *BME280 Combined humidity and pressure sensor*, Bosch, 09 2018.
[13] *DS10314 Rev.8 STM32F411xC STM32F411xE datasheet*, ST Microelectronics, 01 2024.
[14] *NEO-6 u-blox 6 GPS Modules*, u-blox, 12 2011.
[15] O. A. Alduchov and R. E. Eskridge, "Improved magnus form approximation of saturation vapor pressure," 11 1997. [Online]. Available: https://www.osti.gov/biblio/548871
[16] *Crystal Units Surface Mount CX2016SA*, KYOCERA AVX, 03 2024.
[17] *Crystal Units Surface Mount CX3225SA*, KYOCERA AVX, 10 2023.
[18] T. Kim *et al.*, "Reversible instance normalization for accurate time-series forecasting against distribution shift," in *International conference on learning representations*, 2021.