UMusic: In-car Occupancy Sensing via High-resolution UWB Power Delay Profile

Shuai Wang
Southeast University
Nanjing, Jiangsu, China
shuaiwang_iot@seu.edu.cn

Yunze Zeng
Bosch Research
Sunnyvale, CA, USA
zengyunze@gmail.com

Vivek Jain
Bosch Research
Sunnyvale, CA, USA
vivek.jain@us.bosch.com

Parth Pathak
George Mason University
Fairfax, VA, USA
phpathak@gmu.edu

ABSTRACT

Occupancy sensing is essential for vehicle safety and security applications such as seat belt reminders, airbag deployment, intrusion detection, and child-left-behind alerts. This paper presents UMusic, a novel in-car occupancy sensing system that reuses the ultra-wideband (UWB) devices already installed for access control in modern vehicles. However, due to the compact size and metal structure, the in-car environment is full of reflected propagation paths, which cannot be precisely resolved even with UWB's wide-bandwidth feature. To overcome this challenge, UMusic introduces a reflected-path decomposition technique to extract a high-resolution power delay profile (PDP) from the channel impulse response (CIR) provided by commodity UWB devices, enabling precise environmental perception. By comparing PDPs in empty and occupied conditions, UMusic is able to detect the occupancy status in both a sedan and an SUV with multiple passengers across various scenarios. Our results show that UMusic achieves a 90.2% detection rate using a single CIR measurement collected within 50 ms, outperforming the state-of-the-art by 15.7%. When aggregating six consecutive CIR measurements, UMusic reaches 99.4% accuracy, demonstrating its effectiveness for real-world deployment.

CCS CONCEPTS

• **Human-centered computing → Ubiquitous and mobile computing design and evaluation methods**.

KEYWORDS

Power-delay Profile; In-car Occupancy; Wireless Sensing; Ultra-wideband

1 INTRODUCTION

The automotive industry has been undergoing a major transformation over the past century, shifting from engine-centric design to prioritizing passenger experience [14]. Modern cars are no longer just transport vehicles but intelligent ecosystems that enhance safety and comfort for users [4, 81, 86]. For instance, incorporating various sensors, automakers like Ford [13], Honda [19], and Tesla [59] are making significant progress in building advanced collision avoidance, theft protection, and keyless entry solutions [12, 21].

Meanwhile, in-car occupancy sensing is gaining traction in the automotive industry, enhancing interior intelligence [4, 14, 32] for applications such as rear seat belt reminders, optimized airbag deployment [38], vital signs recognition [25, 29, 91], child (or pet) left behind detection [20, 85, 90], and personalized HVAC and stereo settings. Despite the promising applications, most of these techniques assume prior knowledge of the car's occupancy status. In other words, these sensing systems are unable to adaptively customize sensor parameters (e.g., camera rotation) to focus on the passenger for improved sensing performance when the occupancy status is unknown. While weight sensors are commonly used for in-car occupancy detection, they struggle to distinguish heavy luggage from human occupants [32]. Additionally, weight sensors are typically constrained to the front seats, while they are unavailable to the rear seats due to their high cost and installation complexity [35, 85].

To address the need for in-car occupancy detection, this paper introduces UMusic, a precise in-car sensing solution leveraging Ultra-Wideband (UWB) technology for occupancy detection. UMusic utilizes UWB's channel estimation capabilities to analyze the in-car environment. Changes in passenger occupancy alter UWB signal propagation paths, referred to as the power delay profile (PDP), leading to variations in channel information, specifically the accessible channel impulse response (CIR) data [48, 49]. UMusic leverages the existing deployment of UWB technology for access control via digital key services provided by manufacturers such as Volkswagen [40], BMW [2], and Hyundai [5] [1]. Unlike vision-based [12, 22, 37, 51, 62], acoustic-based [90], mmWave-based [16, 25, 58], and WiFi-based [20, 81, 85] approaches, UMusic offers enhanced privacy preservation, cost-effectiveness, high precision and a lightweight design.

However, directly applying CIR data for in-car environment sensing, particularly for occupancy detection, presents unique challenges. The metal structure confines reflected UWB signals within a compact (2m×2m) space, leading to rich reflections with similar propagation path lengths. These reflections are difficult to differentiate in the CIR data due to limited spatial resolution; specifically,

---

[1]These digital key setups typically involve up to eight internal UWB devices [17, 32], with up to four devices installed externally.

two paths need to differ by at least 0.6 m to be separated effectively, as validated in Section 2.3. To address this challenge, deep-learning-based approaches, such as CarOSense [32], are utilized. These methods typically involve complex deep learning models that demand extensive training and testing to generalize across different car models, which poses challenges for their widespread adoption across various vehicle types.

UMusic takes a different approach by employing signal processing techniques combined with a simple classification model to precisely detect in-car occupancy via the following innovative technical highlights: (i), **High-Resolution and Robust PDP Calculation** - UMusic carefully extracts reflected signal paths to compute a high-resolution PDP from CIR data by decomposing in-car reflections. Additionally, we provide a formal proof demonstrating the robustness of our PDP calculation method against errors introduced by UWB hardware imperfections. (ii), **Smart Occupancy Detection** - UMusic effectively utilizes a phenomenon where the presence of a passenger only affects signal propagation paths longer than the line-of-sight (LoS) path (i.e., TX-Passenger-RX), while shorter paths remain unaffected. By excluding these longer paths from the PDP, UMusic reduces the complexity of occupancy detection. With the assistance of a simple machine learning model such as SVM, UMusic detects changes in the LoS paths, enabling accurate occupancy status detection. (iii), **Computational Efficiency** - To optimize performance for onboard computers, UMusic employs a downsampling strategy to lower the computational cost of calculating PDP, cutting it down to a brisk **0.125 ms**. A formal proof also substantiates the method's capability to preserve PDP accuracy efficiently. In addition, we address CIR aliasing and misalignment issues, enhancing the practicality of UMusic. By incorporating these techniques, UMusic achieves a highly accurate, lightweight, cost-effective, and privacy-preserving design, making it suitable for deployment in various car models. We evaluate UMusic's performance on two car models across different scenarios, achieving an overall detection rate of **90.2%** and an accuracy of **99.4%** when aggregating the results of six consecutive estimations. To summarize, the contribution of this paper is threefold:

- We present UMusic, an innovative in-car occupancy sensing system designed to detect passengers' seats accurately. UMusic envisions to utilize the UWB devices available in cars for access control, making it a cost-effective and easy-to-deploy solution.
- UMusic introduces an innovative path decomposition technique to capture the high-resolution PDP of in-car signal propagation, which is then used to detect passenger occupancy. To implement UMusic effectively, we tackle UWB hardware imperfections such as carrier frequency offset, sampling time offset, and random initial phase. Additionally, we reduce computational complexity to ensure compatibility with various onboard computing systems.
- We evaluate UMusic's performance on DW1000 UWB devices under various scenarios. The experimental results demonstrate that UMusic achieves 99.4% accuracy in occupancy detection. At the same time, the simulation shows the high precision of PDP calculation, indicating the potential for UMusic's application in more general scenarios.
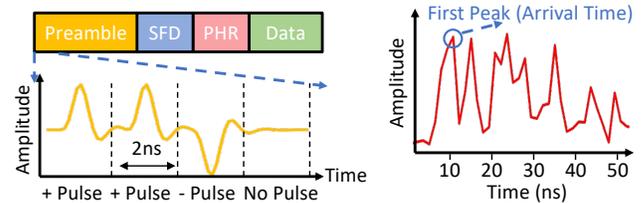
(a) An UWB PHY packet consists of a Preamble, Start Frame Delimiter (SFD), PHY Header (PHR), and Data fields.

(b) The first peak in CIR data is detected as the token of UWB packet arrival.

Figure 1: Illustration of UWB PHY layer and CIR-based packet arrival detection.

2 BACKGROUND AND MOTIVATION

2.1 The Need of In-car Occupancy Sensing

In recent years, there have been significant advancements in in-car human sensing [12, 14, 64, 81], which allows the car to obtain the passengers' positions [25, 32], presence [12, 21], and even vital signs [25, 55]. Of these sensing abilities, occupancy detection has become increasingly important for modern vehicles. For example, by assisting impact sensors, the in-car occupancy detection system can ensure that only specific airbags (i.e., the driver and passenger front airbags and corresponding side airbags) are deployed, while the rest remain deactivated to prevent unnecessary injury [23]. Additionally, the vehicle's interior lighting, audio system, air conditioning, and power accessories can be adapted to enhance the passenger experience. Furthermore, occupancy detection serves as the foundation for many existing in-car sensing systems, including those for detecting vital signs [25, 55], which can be used to enable child left behind and medical emergency detection.

Despite the potential benefits, a deployable in-car occupancy detection system is not there yet. EU NCAP [38] and NHTSA [39] currently only mandate occupancy detection for the driver and front seats, leaving rear seat detection unregulated, while the ability to derive the potential presence of a subject or object inside the car based on pressure will not be rewarded from 2025 onwards. Consequently, weight/pressure sensors are primarily installed in front seats, as equipping all seats with dedicated sensors would increase hardware and installation costs [35, 85]. A car occupancy detection system, in general, should be highly accurate and commercially viable, prompting us to rethink the UWB technologies that are already integrated into the existing in-car systems [2, 5].

2.2 UWB Primer

UWB is a wireless technology widely applied in real-time location systems, featured by NXP Semiconductors [53] and Qorvo [48, 49] radio chips as well as many famous vendors, including Apple, Samsung, and Xiaomi. Its high accuracy in distance ranging is essentially achieved by utilizing UWB CIR data to determine the packet arrival time, which is used to precisely estimate the distance between the transmitter and receiver. The UWB PHY packet begins with the preamble field, consisting of sequences of +/- pulses or no pulse, as illustrated in Figure 1(a). The preamble sequence is designed to maintain a perfect periodic autocorrelation, enabling the receiver to obtain the exact CIR using a correlator [34]. As a result, the receiver marks the timestamp with the arrival time
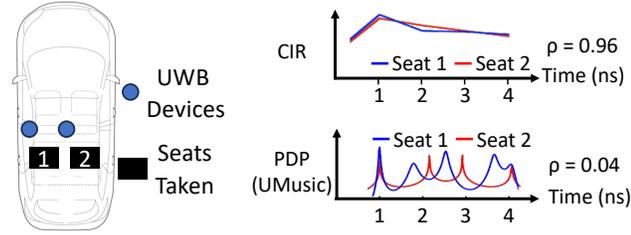
(a) UWB devices are deployed to collect CIR data when a person sits on two different seats. (b) CIR amplitude (upper) collected at two different seats are compared with the corresponding the high-resolution PDP (lower) calculated by UMusic.

Figure 2: Preliminary study in a car to compare the CIR amplitude and PDP calculated by UMusic.

of the first path (the first peak in CIR), as depicted in Figure 1(b), while the full CIR data is temporally saved in the UWB PHY layer. With a pulse duration of 2 ns, which is equivalent to the 60 cm spatial resolution, UWB can support large-area sensing tasks such as people counting [9] and localization [26, 60, 61, 63] in a large room, car localization [82] on the road, and enabling keyless entry for cars [21]. However, directly using CIR amplitude for small-area sensing such as car occupancy detection would result in ambiguity, as demonstrated in our preliminary study in the next section.

2.3 The Limitation of UWB CIR Amplitude

We deploy two UWB devices in the middle of a car, as depicted in Figure 2(a), to collect the CIR data when a person sits on seat 1 and seat 2. As a person sits in different seats, the signal propagation paths affected by the human body should change significantly. The essence of occupancy detection is to capture this change from the power delay profile. However, as shown in Figure 2(b), the CIR data (amplitude) collected on these two seats are quite similar (with a correlation ρ of 0.96). Such a high correlation would eventually result in the ambiguity of directly applying CIR to detect the occupied seats. This happens because the signal propagation change caused by these two occupancy statuses is much less than the spatial resolution (60 cm). Since the CIR amplitude couldn't fully capture any signal propagation change within 60 cm, it is challenging to distinguish the two adjacent seats (20 to 30 cm away) occupancy.

Previous work on using UWB for small-area sensing has either relied on heavy machine-learning models [32, 89, 92] or complex hardware setups [3, 71, 88]. However, these designs are impractical for low-cost UWB radio chips and onboard computers with limited computational resources. In contrast, UMusic is designed to work with commodity UWB devices via a novel signal-processing technique to extract a high-resolution PDP that can perceive small environment changes caused by different occupancy statuses. As illustrated in Figure 2(b), the PDP calculated by UMusic reflects signal propagation in finer granularity, with the common first peak corresponding to the Tx-Rx path, while other peaks containing the Tx-Body-Rx path are different for the two seats. With the help of a lightweight classification model, UMusic achieves a detection rate of 99.4%, aggregated from six consecutive CIR measurements. In the next section, we provide an overview of UMusic, followed by a detailed design of each technical highlight.
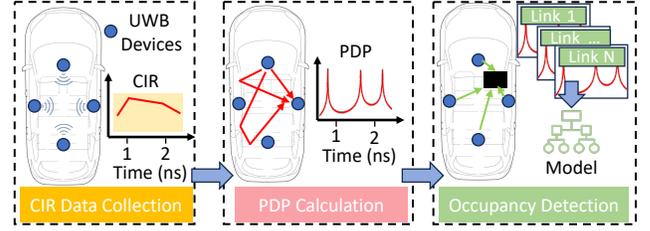
Figure 3: UMusic operates in three steps, from collecting CIR data using UWB devices deployed in the car to obtaining the high-resolution PDP, from which the occupancy is detected.

3 DESIGN OVERVIEW

As depicted in Figure 3, UMusic comprises three steps: (i) CIR data is collected from multiple links of transmitters and receivers installed in a vehicle. The multiple Tx-Rx links allow us to obtain the PDP from various angles, providing a better perception of the signal propagation changes caused by a human. (ii) UMusic calculates the high-resolution PDP from the collected CIR data. Finally, the PDPs obtained from multiple Tx-Rx links are fed into a simple classification model in Step (iii) to detect the car occupancy status.

The resolution of CIR amplitude is limited by the UWB PHY layer design, posing significant challenges for PDP calculation. UMusic overcomes this limitation and achieves high-resolution PDP calculation and accurate car occupancy detection by fully investigating the CIR data's amplitude and phase information. However, due to hardware imperfections, the phase information in CIR is highly biased, leading to more practical issues in the design of UMusic. The following section presents our solutions for these challenges.

4 MAIN DESIGN

As the core of UMusic, decomposing the reflected path in high-resolution PDP calculation is demonstrated first in this section, followed by car occupancy detection.

4.1 High-resolution Power Delay Profile in UWB
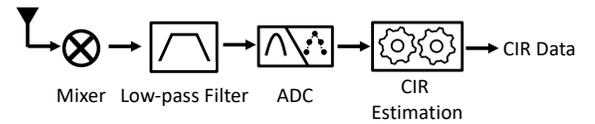
Figure 4: UWB PHY layer passes the received signal to Mixer, Low-pass filter, ADC, and CIR Estimation, to extract CIR.

To illustrate the root cause of the ambiguity in CIR amplitude, we first formulate the CIR estimation process in the UWB PHY layer before demonstrating the detailed design for high-resolution PDP calculation. Let $x(t)$ denote the transmitted signal of the preamble field in the UWB packet. After the UWB signal is transmitted, it traverses through N signal propagation paths, resulting in copies with a delay of τ_i for the i-th path. These delayed copies arrive at the receiver side consecutively, yielding the received signal $y(t)$:

$$y(t) = \sum_{i=0}^{N-1} a_i\, e^{-j2\pi f_c \tau_i} x(t - \tau_i) \tag{1}$$

where f_c is the center frequency of the UWB signal, and a_i is the attenuation of the i-th path. In addition, the corresponding power delay profile could be formulated as follows:
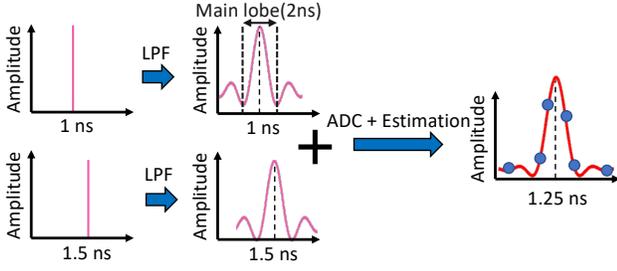
Figure 5: Two paths with the delay of 1 ns and 1.5 ns are added up to be the CIR data peaking at 1.25 ns.

$$h(t) = \sum_{i=0}^{N-1} a_i\, e^{-j2\pi f_c \tau_i} \delta(t - \tau_i) \tag{2}$$

where $\delta(t)$ is a Dirac's delta function [41]. The received signal $y(t)$ is fed into the UWB PHY layer to obtain the CIR data as an estimation of $h(t)$, as illustrated in Figure 4.

Specifically, upon the arrival of the UWB signal, the Mixer first performs a passband to baseband conversion on the signal to obtain the baseband signal $y(t)$. Subsequently, the received signal is passed through a Low-pass Filter to eliminate unwanted parts and retain the middle B Hz bandwidth (e.g, 500MHz), which is achieved mathematically by convolving $y(t)$ with the sinc-shaped filter, represented by $\mathrm{sinc}(Bt) = \frac{\sin(\pi Bt)}{\pi Bt}$ [41]. The resulting signal is then sampled by the Analog-to-digital Converter (ADC) at every T_s (e.g., 1 ns) interval to generate time-domain samples. Finally, the CIR Estimation applies a correlator to extract the CIR Data, which is denoted by $h[n]$:

$$h[n] = \sum_{i=0}^{N-1} a_i\, e^{-j2\pi f_c \tau_i} \mathrm{sinc}\left(n - \frac{\tau_i}{T_s}\right) \tag{3}$$

By comparing $h[n]$ in Equation 3 and $h(t)$ in Equation 2, the precision of using CIR amplitude as PDP is determined by the sinc filter. To demonstrate these effects thoroughly, we provide an example with two signal propagation paths.

Figure 5 illustrates an example where the channel includes two paths with propagation delays of 1 ns and 1.5 ns, respectively. When two copies of the transmitted signal arrive through these paths, they are shaped by the sinc filter with a main lobe of 2 ns, resulting in the corresponding two peaks merging into a single high peak at 1.25 ns. After ADC sampling, the resulting samples are further distorted, and the CIR amplitude does not accurately capture the two paths of 1 ns and 1.5 ns. This ambiguity in the CIR data can mislead car occupancy detection. To address this ambiguity, Caraokey [21] attempts to improve the precision of CIR amplitude through zero-padding-based upsampling. However, this approach only smooths the CIR amplitude, and the ambiguity remains unresolved. In the following section, we will demonstrate how UMusic overcomes the influence of LPF and achieves a high-resolution PDP.

4.1.1 Reverting Low-pass Filter. In order to undo the LPF's sinc shaping, the time-domain CIR $h[n]$ is transformed into the frequency domain, resulting in the channel frequency response (CFR):

$$H[k] = \sum_{i=0}^{N-1} a_i\, e^{-j2\pi(f_0 + (k-1)f_\Delta)\tau_i} \tag{4}$$

where f_0 represents the center frequency of the leftmost frequency bins, and f_Δ denotes the channel spacing, which is typically 5 MHz in UWB. As the sinc function in the frequency domain is equivalent to a rectangular function [41], the CFR in Equation 4 is free of LPF's influence. While this calculation is simple, it is a crucial step for reverting the low-pass filter to obtain the reflected paths.

4.1.2 Reflected Paths Separation. However, transforming the CIR data into the frequency domain incorporates all path information into each frequency bin, making it challenging to distinguish individual paths. To clarify this problem, let us first establish the formulation. Specifically, we represent the phase increment between two consecutive frequency bins for the i-th path with $\Omega_i = e^{-j2\pi f_\Delta \tau_i}$. Similarly, we use $\gamma_i = a_i e^{-j2\pi f_0 \tau_i}$ to denote the complex value of the i-th path on the first (leftmost) frequency bin. Using these notations, the received CFR of the N paths are expressed as follows:

$$\underbrace{\begin{bmatrix} H[1] \\ H[2] \\ \vdots \\ H[M] \end{bmatrix}}_{\mathbf{H}(M\times 1)} = \underbrace{\begin{bmatrix} 1 & 1 & & 1 \\ \Omega_1 & \Omega_2 & & \Omega_N \\ \vdots & \vdots & \cdots & \vdots \\ \Omega_1^{M-1} & \Omega_2^{M-1} & & \Omega_N^{M-1} \end{bmatrix}}_{\Omega(M\times N)} \underbrace{\begin{bmatrix} \gamma_1 \\ \gamma_2 \\ \vdots \\ \gamma_N \end{bmatrix}}_{\Gamma(1\times N)} \tag{5}$$

where M is the number of frequency bins included in the PDP calculation; typically 100 in UWB. Essentially, the process of computing PDP from CIR data involves solving Equation 5 to derive the elements of matrix Ω. However, due to the presence of two unknowns (Ω and Γ) on the right-hand side and only one known (CFR vector $\mathbf{H}$) on the left-hand side, this equation cannot be straightforwardly solved using standard linear algebra techniques.

We note that the matrix Ω exhibits the Vandermonde property, which motivates us to utilize the MUSIC algorithm [52] to solve Equation 5. In accordance with the MUSIC algorithm's convention, we refer to the matrix Ω as the steering matrix, where each column is known as a steering vector. Originally developed for determining the angle-of-arrival (AoA) of each incident signal in a uniform linear antenna array, the MUSIC algorithm accepts the known left-hand side matrix (corresponding to $\mathbf{H}$ in Equation 5) as input and generates an estimate for every element in the steering matrix of Vandermonde shape (which corresponds to $\mathbf{H}$ in Equation 5), without requiring knowledge of Γ. It is worth pointing out that the MUSIC algorithm assumes uncorrelated incident signals and requires a sufficient number of signals to be collected from the antenna array. Consequently, to apply the MUSIC algorithm, $\mathbf{H}$ and Γ must be matrices with a rank greater than the number of reflected paths N, as opposed to vectors in Equation 5. The next section illustrates how to transform $\mathbf{H}$ and Γ into the required matrices, while preserving the same steering elements, to prepare them to be solved by the MUSIC algorithm.

4.1.3 CFR Transformation. The transformation of the CFR vector is motivated by Spatial Smoothing [54], which rearranges the CFR vector $\mathbf{H}$ into a full-rank matrix without altering any of the elements in the steering matrix. The essential idea behind this procedure is that if we could identify several subsets of $\mathbf{H}$ with the same steering elements as the initial $\mathbf{H}$, these subsets would be combined to form a full-rank matrix. The details are explained as follows:

$$H[1] \quad H[2] \quad \ldots \quad H[p] \quad H[p+1] \quad \ldots \quad H[q] \quad \ldots \quad H[q+p-1] \quad \ldots \quad H[M]$$

$$
\begin{bmatrix} H[1] & H[2] & & H[q] \\ H[2] & H[3] & & H[q+1] \\ \vdots & \vdots & \cdots & \vdots \\ H[p] & H[p+1] & & H[q+p-1] \end{bmatrix} = \begin{bmatrix} 1 & 1 & & 1 \\ \Omega_1 & \Omega_2 & & \Omega_N \\ \vdots & \vdots & \cdots & \vdots \\ \Omega_1^{p-1} & \Omega_2^{p-1} & & \Omega_N^{p-1} \end{bmatrix} \begin{bmatrix} \gamma_1 & \Omega_1\gamma_1 & & \Omega_1^{q-1}\gamma_1 \\ \gamma_2 & \Omega_2\gamma_2 & & \Omega_2^{q-1}\gamma_2 \\ \vdots & \vdots & \cdots & \vdots \\ \gamma_N & \Omega_N\gamma_N & & \Omega_N^{q-1}\gamma_N \end{bmatrix}
$$

$$\underbrace{}_{\mathbf{H'}(p \times q)} \quad \underbrace{}_{\mathbf{\Omega'}(p \times N)} \quad \underbrace{}_{\mathbf{\Gamma'}(N \times q)}$$
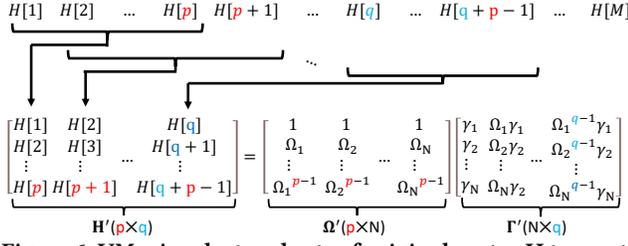
Figure 6: UMusic selects subsets of original vector H to create a new matrix H′.

The selection of subsets of p elements from **H** to form the columns of the new matrix **H′** is depicted in Figure 6. The new steering matrix, $\mathbf{\Omega'}$, is identical to the original $\mathbf{\Omega}$. Since **H′** is the product of the full-rank matrix $\mathbf{\Omega'}$ and $\mathbf{\Gamma'}$, which is a scaled version of the Vandermonde matrix and also full-rank, UMusic successfully transforms the CFR vector **H** into a full-rank matrix **H′**. The size of the matrix $\mathbf{\Omega'}$ is $p \times N$, while the size of $\mathbf{\Gamma'}$ is $N \times q$. As both matrices are full-rank, the rank of $\mathbf{\Gamma'}$ is $\min(q, N)$, and the rank of **H′** is $\min(p, q, N)$. To meet the rank requirement for **H′** and $\mathbf{\Gamma'}$, which must exceed the number of reflected paths N, we select values of p and q that are greater than N. This ensures that the requirement for using the MUSIC algorithm to obtain the PDP is satisfied. We can then input the matrix **H′** into the MUSIC algorithm to obtain the high-resolution PDP, consisting of τ_1 to τ_N.

In UWB, the number of frequency bins, M, is typically 100, and the number of reflected paths within the $2m \times 2m$ in-car area is limited. This implies that there are many choices for the values of p and q. However, if we carelessly select the combination of p and q, it could lead to poor PDP performance. An example to illustrate this is provided next.
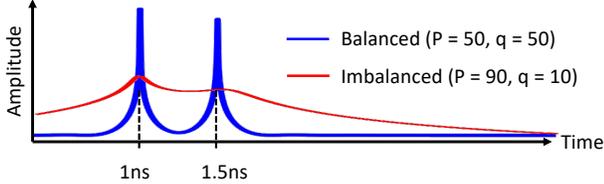
Figure 7: PDP calculation result using balanced and imbalanced settings for the matrix $\mathbf{\Omega'}$ and $\mathbf{\Gamma'}$.

Consider the two paths shown in Figure 5 as an example. The resulting PDP calculated by ($p = 50, q = 50$) has more significant peaks than the result obtained with ($p = 90, q = 10$), as illustrated in Figure 7. This is because the corresponding steering vectors $\overrightarrow{\Omega_1} = [1, \Omega_1, \ldots, \Omega_1^{p-1}]^\top$ and $\overrightarrow{\Omega_2} = [1, \Omega_2, \ldots, \Omega_2^{p-1}]^\top$ are less distinguishable under ($p = 90, q = 10$) than under ($p = 50, q = 50$). Specifically, the maximum phase change for each steering vector is given by Ω_1^{p-1} and Ω_2^{p-1}, which are $\frac{(p-1)\pi}{100}$ and $\frac{1.5(p-1)\pi}{100}$, respectively. When $p = 10$, the difference between these two phases is only 0.095π, which is insufficient to distinguish between the two paths. In a practical low SNR environment, the results could be even more degraded. Since the values of p and q are constrained by the number of frequency bins, M, such that $p + q \leq M$, UMusic must carefully select p and q to maximize the phase change over the steering vectors. Therefore, by selecting $p = q = M/2$, UMusic maximizes the phase change over the steering vectors, which leads to the maximum achievable SNR.
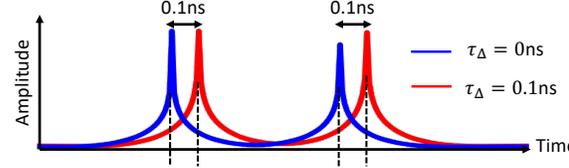
Figure 8: The examples of PDP estimation with a sampling time offset (red) and without a sampling time offset (blue).

4.1.4 Immunity to Hardware Imperfections. As the CIR data is obtained from the UWB PHY layer, any hardware imperfections in the components of the PHY layer may result in an inaccurate matrix **H′** and subsequently affect the PDP calculation. In this section, we will show that our PDP calculation is robust and can be immune to the hardware imperfections.

Essentially, hardware imperfections can introduce three major errors in the CIR data, including: (i) carrier frequency offset (CFO), denoted as f_{CFO}, which arises due to unsynchronized oscillators in the UWB transmitter and receiver, causing a slight mismatch between the mixer at the receiver side and the center frequency of the transmitted UWB signal; (ii) sampling time offset τ_Δ, which is caused by unsynchronized analog-to-digital converters (ADCs) that sample the signal with a random time shift; and (iii) initial phase offset θ_Δ, which is an inherent and unknown phase value imposed on RF devices when they are powered on. In the presence of these imperfections, the received CFR is formulated as $\hat{H}[k]$:

$$
\begin{aligned}
\hat{H}[k] &= \sum_{i=0}^{N-1} a_i\, e^{-j2\pi(f_{CFO}+f_0+(k-1)f_\Delta)(\tau_i+\tau_\Delta)+\theta_\Delta} \\
&= \sum_{i=0}^{N-1} \hat{\Omega}_i^{\,k-1}\hat{\gamma}_i
\end{aligned}
\tag{6}
$$

where $\hat{\Omega}_i = e^{-j2\pi f_\Delta(\tau_i+\tau_\Delta)}$ and $\hat{\gamma}_i = a_i e^{-j2\pi(f_{CFO}+f_0)(\tau_i+\tau_\Delta)+\theta_\Delta}$. It is important to note that the effects of CFO and random initial phase are confined to $\hat{\gamma}_i$, which implies that these hardware imperfections do not affect the steering matrix. Additionally, since the MUSIC algorithm does not require knowledge of $\hat{\gamma}_i$ to calculate the steering matrix, any variations in $\hat{\gamma}_i$ will not alter the PDP calculation result. Therefore, our high-resolution PDP calculation is immune to the effects of CFO and random initial phase.

The sampling time offset affects all estimated paths simultaneously by introducing a delay of τ_Δ, which results in a shift of all estimated peaks by the same amount. As shown in Figure 8, all peaks have a delay of 0.1 ns compared to the case without sampling offset. The relative distance between two peaks (paths) remains unchanged, irrespective of the sampling offset, which is critical for aligning all calculated PDPs, as described in Section 5.3. With this analysis, it can be concluded that our high-resolution PDP is immune to the three hardware imperfections mentioned above. This immunity is a unique advantage of our approach compared to existing wireless sensing works, such as SpotFi [24] and Seirious [27], which are significantly impacted by these imperfections and require additional designs to compensate for their effects. Moreover, unlike SpotFi, which relies on a multi-antenna setup in WiFi to decompose reflected paths, UMusic requires only a single antenna configuration in UWB, thanks to its novel formulation for UWB signal decomposition and optimal solution derivation. In the next
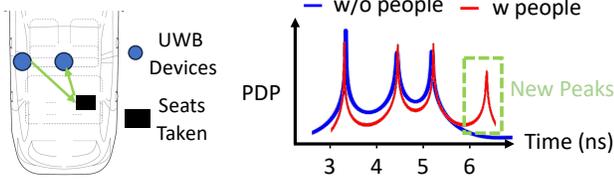
(a) UWB devices are deployed to collect CIR data when a person sits on the right back seat.

(b) Comparison between the PDP of with people and without people shows the human body creates a new reflection (peak marked in the green box).

Figure 9: Example of how human occupancy changes the in-car signal propagation.

section, we will demonstrate how UMusic detects car occupancy by utilizing the high-resolution PDP.

4.2 Car Occupancy Detection using PDP

The detection of car occupancy primarily relies on our observation that the human body only affects signal propagation paths that are longer than the LoS path (TX-body-RX), while the shorter paths remain unaffected. To illustrate this observation, we compare the PDP obtained under two conditions: an empty car (without people occupancy) and a car with a person sitting in the right back seat, as shown in Figure 9. The results in Figure 9(b) demonstrate the newly created reflection caused by human occupancy. Specifically, the newly created reflection has a delay of 6.4 ns, which is 3.1 ns later than the first path (3.3 ns), while the remaining three peaks remain unchanged. This delay of 3.1 ns corresponds to a distance of 0.93 m, which is the additional length of the Tx-body-Rx path compared to the Tx-Rx path (1 m), thereby confirming the validity of our observation.

Using this observation and the high-resolution PDP, car occupancy detection can be achieved by comparing the PDP collected without people to the PDP obtained under the current status. As the in-car layout is limited to a $2m \times 2m$ area, we only compare the PDP within the first 4 meters for occupancy detection, while PDPs longer than that are disregarded as they correspond to non-Tx-body-Rx paths. To enhance occupancy detection, UMusic utilizes the deployment of multiple UWB devices and simple classification models. Our evaluation shows that, with the support of high-resolution PDP, even traditional classification models like SVM can reach an accuracy of 90.2% for in-car occupancy detection.

5 EFFICIENCY ENHANCEMENT

UMusic's efficiency is further enhanced by three additional features.

5.1 Computational Cost Optimization

This section reduces computational costs to facilitate UMusic's integration into onboard systems without overburdening resources already running multiple applications. As the solver for high-resolution PDP calculation, the MUSIC algorithm takes $O(p^3)$ time complexity, which is dominated by the heavy eigenstructure decomposition of the covariance matrix of $\mathbf{H}'$ (size of $p \times q$) [18]. However, since p and q have already been optimized, directly reducing these parameters would lead to a decrease in the precision of the PDP.

Our approach to achieving complexity reduction without compromising PDP precision is depicted in Figure 10. Here, we reduce the size of $\mathbf{H}'$ by a factor of D by downsampling the matrix $\mathbf{H}'$,

Figure 10: Downsampling the matrix $\mathbf{H}'$ by D to reduce the computational cost.

resulting in a complexity reduction of D^3. For ease of understanding, we assume that p and q are multiples of D. By evenly selecting one CFR element from every D elements, we obtain a downsampled CFR matrix, referred to as $\mathbf{H}^D$, which can be decomposed into two matrices, namely Ω^D and Γ^D. Although the i-th downsampled steering vector $\overrightarrow{\Omega_i^D} = [1, \Omega_i^D, \ldots, \Omega_i^{p-D}]^\top$ has only $\frac{1}{D}$ of the elements compared to the steering vector $\overrightarrow{\Omega_i} = [1, \Omega_i, \ldots, \Omega_i^{p-1}]^\top$, the PDP calculated from this downsampled matrix remains accurate.

The high precision is ensured by the substantial phase change over the downsampled steering vector. The phase change over the downsampled steering vector $\overrightarrow{\Omega_i^D}$ is given by $2\pi f_\Delta (p - D)\tau_i$, which is $2\pi f_\Delta (D-1)\tau_i$ less than the phase change over the original steering vector. The lost, $2\pi f_\Delta (D-1)\tau_i$, is negligible, and thus downsampling does not alter the correlation between the steering vectors of different paths.
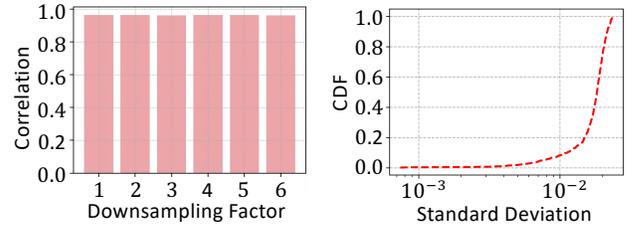
(a) The correlation between steering vector of 1 ns and 1.5 ns delay over the downsampling factor.

(b) The CDF of the STD of the correlation between 1000 simulated steering vectors, aggregated from D = 1 to 6 .

Figure 11: The correlation of two steering vectors maintains the consistency over the downsampling factor D.

This is validated in Figure 11(a), where the correlation between $\overrightarrow{\Omega_1^D}$ and $\overrightarrow{\Omega_2^D}$ is compared at different downsampling factors, even a six-fold reduction. The results indicate that the correlation between $\overrightarrow{\Omega_1^D}$ and $\overrightarrow{\Omega_2^D}$ remains relatively close, even after downsampling. To provide a more comprehensive evaluation, we simulate the steering vectors of 1000 different paths and calculate their correlation for various downsampling factors. The results, shown in Figure 11(b), indicate that the accuracy of the PDP calculation remains unaffected and consistent across different D values. Since the MUSIC algorithm primarily relies on the correlation between different paths, preserving this correlation is crucial for downsampling to reduce computational costs without losing PDP precision.

5.2 Aliasing Avoidance

Despite the effectiveness of computational cost reduction, the side effect of downsampling is aliasing, where two different path yields the same steering vector. This is formally defined as $\overrightarrow{\Omega_1^D} = \overrightarrow{\Omega_2^D}$, while $\tau_1 \neq \tau_2$. Then, the two aliased paths follow:

$$2\pi D f_\Delta \tau_1 = 2\pi D f_\Delta \tau_2 \mod 2\pi$$

$$\Longleftrightarrow \quad \tau_1 = \tau_2 - \frac{r}{D f_\Delta}, r \geq 1 \tag{7}$$
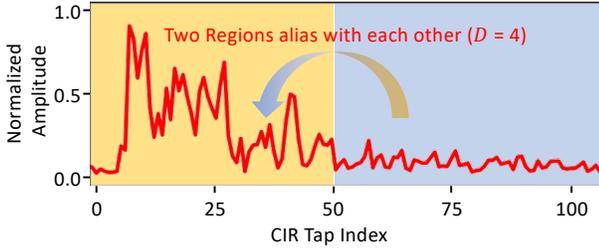
Figure 12: The amplitude of CIR data collected from an in-car experiment demonstrates the long tail trend for 100 taps (or equivalently 100 ns).

where r is an integer. Aliasing causes erroneous results in high-resolution PDP calculation because the longer path (i.e., τ_2) confuses the computation of the short path (i.e., τ_1) if they satisfy the relationship specified in Equation 7. In fact, this relationship is easy to meet in the in-car environment.

Equation 7 implies that aliasing happens $\tau_1 = \tau_2 - 50$ ns, when $D = 4$, as illustrated in Figure 12, where the aliasing problem between the two regions with 50 ns (taps) gap. Since in-car signal propagation has a fairly long tail due to the signal frequently bouncing in compact and metal car structures, the aliasing issue is inevitable if using downsampling.
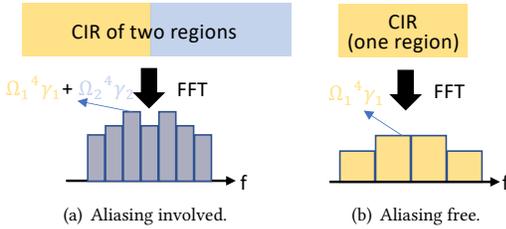
Figure 13: Illustration of aliasing caused by involving two regions in FFT calculation and aliasing avoidance via excluding the second region.

To overcome the aliasing issue, we exclude the regions that could alias the first eight CIR taps. This protects the PDP calculation for the first 2.4 meters environment because this range is mainly leveraged for in-car occupancy detection as illustrated in Section 4.2. As illustrated in Figure 13, by excluding the second region, FFT results only contain the paths within the first region, thereby free from aliasing. This also brings another benefit: involving less number of CIR taps into FFT calculation naturally expands the channel spacing, i.e., f_Δ, which is equivalent to downsampling the CFR matrix.

5.3 PDP Synchronization

Affected by the sampling offset, the calculated PDP results suffer from severe misalignment, as depicted in Figure 14(a). As of 1ns sampling offset, the corresponding misalignment (30 cm) is significant enough to cause the wrong comparison of PDP, even misleading the occupancy detection result. To synchronize the PDP calculated from different UWB packets, we leverage the insight in Section 4.1.4, where the relative distance between the peaks in PDP is immune to the sampling offset. Specifically, the first peak in the PDP result corresponds to the Tx-Rx path, which is the shortest one compared with all other paths. This path length is fixed since
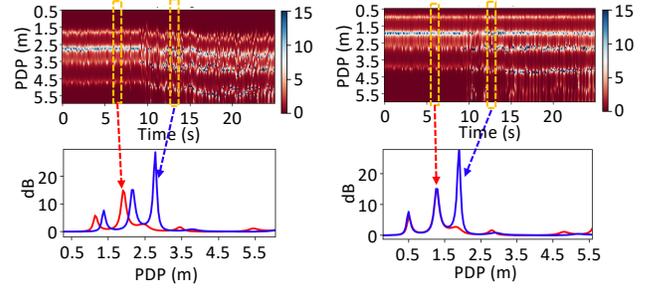
(a) The spectrum of unsynchronized PDP results, where two PDPs are selected to show the misalignment.

(b) The spectrum of synchronized PDP results, where two PDPs are selected to show the alignment.

Figure 14: The PDP results with/without synchronization.

the Tx and Rx are tightly attached to the car. Therefore, by aligning the first peak in the PDP results, all PDPs are synchronized. As Figure 14(b) shows, the second peaks in two PDPs are also aligned after the synchronization of the first peak. This concludes all designs of UMusic. Next, we will demonstrate the evaluation of UMusic.

6 EVALUATION

6.1 Implementation

Figure 15: UMusic leverages the UWB devices deployed to the eight locations inside a sedan parked in a garage.

System Implementation: We implement UMusic in Python codes (i), collect CIR data from Radino DW1000 [2] modules [50], (ii), calculate high-resolution PDP. (iii), recognize the car occupancy status via representative classification models. Moreover, we implement CarOSense as the State-of-the-Art (SoTA). Our approach for processing the CFR vector $\mathbf{H}$ is based on the PyArgus tool [45]. We select an SVM model as an occupancy detection model, implemented by directly calling svm.SVC() by default provided by scikit-learn [44]. Moreover, UMusic is implemented on a Mac M1 computer configured to use a single CPU and a single thread, without relying on GPU acceleration or any specific hardware needs from the Mac M1.

SoTA Approach: CaroSense [32] feeds the raw CIR data collected from eight UWB sensors deployed in a car into a deep learning model, which incorporates MIMO and masking techniques to determine occupancy status.

Data Collection: Our evaluations include comparing UMusic with SoTA system, CarOSense [32] utilizing the same setup and dataset, for fairness. As Figure 15 depicts, eight UWB devices are positioned in a **sedan** at the following locations: 1 (front-left ceiling), 2 (rear-view mirror), 3 (front-right ceiling), 4 (center panel), 5 (back-center ceiling), 6 (back-left ceiling), 7 (back-right ceiling), and 8

---

[2]DW1000 is chosen to demonstrate the concept as it is commercially available [11] and the UMusic is easily portable to any UWB device as all of them support CIR extraction.
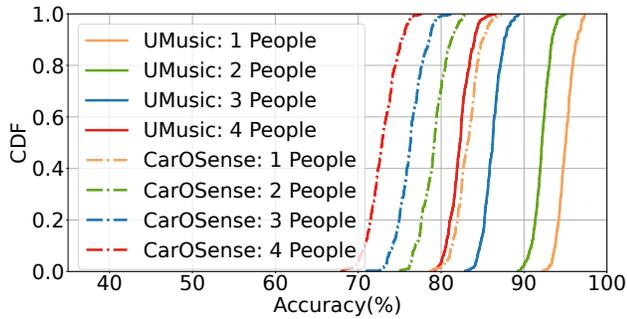
Figure 16: Comparison of UMusic and CarOSense in detecting various numbers of people in a car.

(trunk). We also evaluate UMusic in an **SUV** with only four UWB devices deployed to showcase its generalizability with fewer UWB devices in Section 6.6. We employ nine volunteers with different bio-metrics (height: [165cm, 182cm]; weight: [125lb, 185lb]) [3]. We evaluate UMusic and SoTA for one to four people occupancy detection, under various scenarios, including **car status** (stationary or driving), **car models** (sedan or SUV), and **out-car environment** (indoor or outdoor). Volunteers are instructed to sit in different positions within the car and are free to move their bodies, hands, and legs, allowing them to adopt relaxed postures, such as using smartphones or talking with other passengers. For example, when evaluating UMusic's performance for detecting 2 people, two volunteers occupy all possible combinations of two seats. The CIR data is collected in a round-robin manner, where each UWB device takes turns transmitting UWB packets, while the other devices collect CIR data. During each round of data collection, data is collected for 10 minutes, with 20 CIR sample data collected every second, resulting in 12K CIR sample data per round. The dataset is divided into training, validation and testing set in a 7:2:1 ratio. **All experiments are approved by our institution's Institutional Review Board (IRB).**

Evaluation Metric: We utilize accuracy to evaluate the performance of in-car occupancy detection. Since volunteers are instructed to occupy all possible seat combinations, we also measure accuracy per seat to evaluate performance in detail for each individual seat.

6.2 Overall Performance

This experiment evaluates UMusic and CarOSense with the car parked in a garage, while no other objects are outside the car. As illustrated in Figure 16, we compare the accuracy of UMusic and CarOSense across different configurations for the number of people. The median accuracy of UMusic for 1 to 4 people is 97.2%, 93.3%, 87.2%, and 83.1%, respectively, representing improvements of 16.8%, 17.9%, 14.3%, and 13.7% compared to CarOSense, which achieves median accuracies of 83.2%, 79.1%, 76.3%, and 73.1%, respectively. This result demonstrates that our high-resolution PDP effectively captures the significant changes in signal propagation due to human occupancy, allowing traditional classification models like SVM to achieve an overall accuracy of 90.2%, outperforming CarOSense by 15.7%. Moreover, the detailed FP and FN of UMusic and CarOSense

---

Number of People	CarOSense		UMusic	
	FP	FN	FP	FN
1 People	8.7%	8.1%	1.5%	1.3%
2 People	11.5%	8.4%	3.2%	3.5%
3 People	10.9%	12.8%	5.8%	7.0%
4 People	12.1%	14.8%	9.3%	7.6%

Table 1: Comparison of False Positives (FP) and False Negatives (FN) between CarOSense and UMusic.

are provided in Table 1, where UMusic achieves a lower error rate with respect to the number of people. Notably, this accuracy could be further improved by aggregating results from the CIR collected at multiple time slots, as discussed in Section 6.5. Moreover, although the posture of lying in the backseat is not included in the experiments, this posture should still signal to UMusic that both backseats are occupied. To illustrate the detailed accuracy per seat, we provide a breakdown of the overall evaluation next.
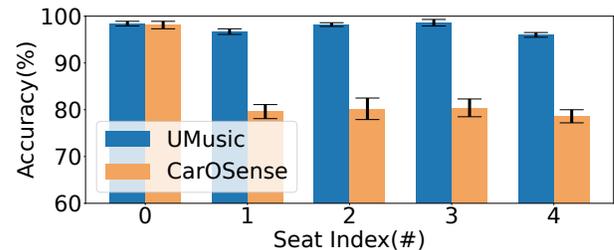
Figure 17: Accuracy under single-person scenario.

6.2.1 Single-person Detection. Detecting a single person is essential for identifying when a person is left in the car. This experiment demonstrates the accuracy of UMusic and CarOSense in detecting a single person per seat in the car. The results are shown in Figure 17, where seat index 0 indicates nobody is in the car. The accuracy varies across different seats but remains above 95%, with the highest accuracy of 98.6% achieved at Seat #3; this is a 28.2% improvement compared to CarOSense's 80.4%. Similarly, UMusic's accuracy is higher than CarOSense across all seats.
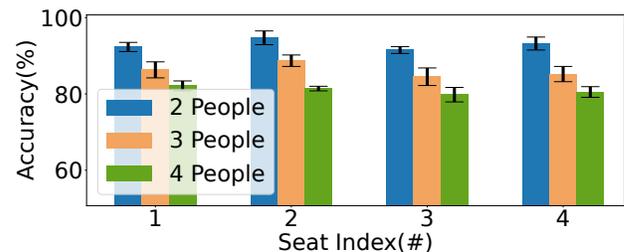
Figure 18: Accuracy for multiple-people detection.

6.2.2 Mutiple-people Detection. As shown in Figure 18, when the number of occupants in the vehicle gradually increases, the accuracy of UMusic experiences a decline. For Seat #1, the accuracy is 92.3% when there are two people in the vehicle, which drops to 86.3% with three occupants, and further decreases to 82.1% when the number of occupants increases to four. The results for other seats are similar to those for Seat #1, as the reflected paths in the confined space of the vehicle become more complex with an increasing number of occupants, leading to deviations in PDP estimation.

---

[3]We do not evaluate UMusic for detecting small children due to the necessity of thorough IRB processing; however, the recently available standardized in-cabin child substitutes [1] could be useful for such evaluations in the future.
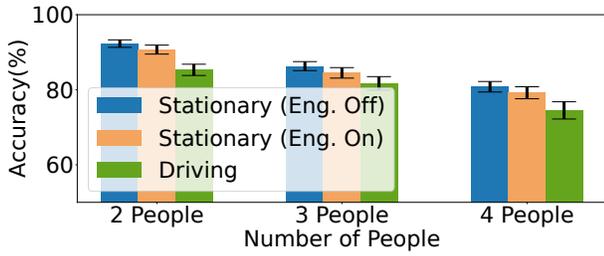
Figure 19: Accuracy of UMusic under stationary (engine on/off) and driving scenarios.

6.3 Stationary vs Driving

To evaluate the robustness of UMusic, we conduct experiments in driving scenarios, where one person occupies the driver's seat while others take turns sitting in the other three seats. The results in Figure 19 indicate that driving and stationary (engine on) have a slight impact on occupancy detection accuracy compared to when the car is stationary and the engine is off. This is likely due to the random shakes experienced when during these two cases, which can affect the collection of CIR data since the attachment to in-car objects may not be entirely stable.
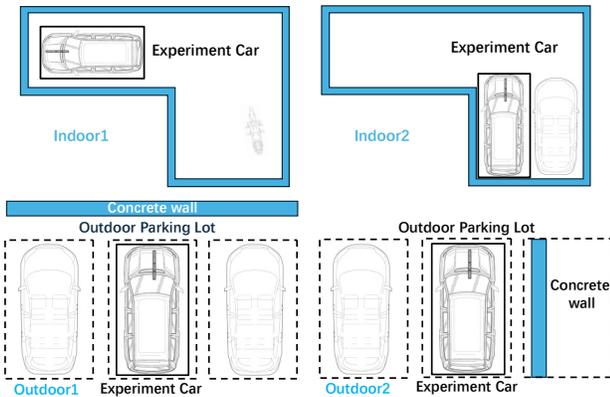
Figure 20: Different out-car environments.

6.4 Impact of Out-car Environments

The metal structure of a car confines wireless signals, preventing the weak in-car RF signal from leaking out of a car and the out-car signals from penetrating into the car. To apply UMusic to the practical scenarios, we evaluate the performance under different out-car environments, as depicted in Figure 20. Figure 21 illustrates the performance of UMusic under different out-car environments. When there is only one occupant in the car, UMusic can accurately detect the occupant, with detection accuracy in different scenarios all above 94%. As the out-car environment changes, there is no significant change in detection accuracy, indicating that the out-car environment has a negligible impact on UMusic for in-car occupancy detection.

6.5 Aggregated Performance

Although the CIR data is collected every 50 ms using the Radino UWB DW1000 module and UMusic can provide results within that interval, users may only require the in-car occupancy system to
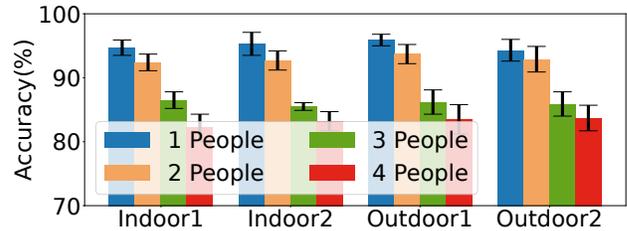
Figure 21: Accuracy under different out-car environments.
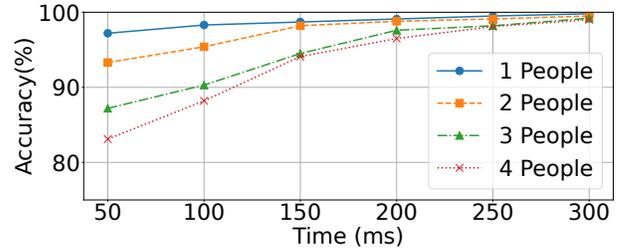
Figure 22: UMusic's accuracy when aggregating multiple detection results, each obtained within 50 ms.

deliver results once per second. This allows us to aggregate detection results over multiple CIR data samples to improve overall performance during that period. In this experiment, we combine the results from multiple detections and use a majority vote to estimate the in-car occupancy more accurately. As shown in Figure 22, after aggregating two to six occupancy detection results, UMusic achieves the accuracy of 93.5%, 96.7%, 98%, 98.4%, and 99.4%, respectively. This also shows that UMusic can achieve over 98% occupancy detection accuracy with just three aggregated results.

6.6 Impact of Different Car Models and UWB Devices Deployment
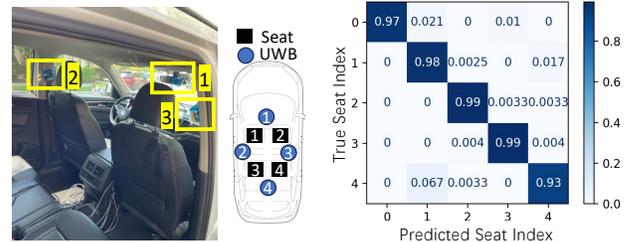
(a) Deployment in an SUV comprises four UWB devices attached to the rearview mirror, left and right pillar, and the back seat.

(b) UMusic performance in detecting occupancy for each seat, with seat index #0 indicating that the SUV is unoccupied.

Figure 23: UMusic performance in an SUV.

To verify the versatility of UMusic, we evaluate occupancy detection in each seat in an SUV with only four UWB devices deployed, as depicted in Figure 23(a). Four UWB devices provide CIR data of 12 links, from which UMusic estimates the occupancy status. Figure 23 shows the occupancy detection results of aggregating six consecutive estimations. Although only four devices are employed in this experiment, the accuracy of occupancy detection remains high in the SUV, all above 93%, showing the effectiveness of UMusic.

Shuai Wang, Yunze Zeng, Vivek Jain, and Parth Pathak
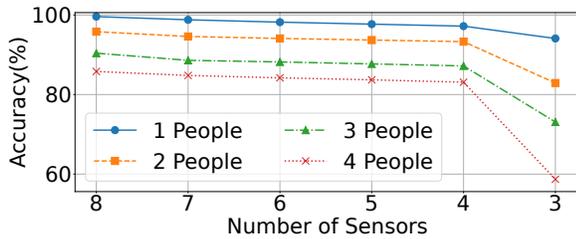
Figure 24: Impact of the number of sensors.

6.7 Impact of the Number of UWB Sensors

To evaluate the impact of the number of sensors on UMusic's performance, we gradually reduce the number of sensors deployed in the vehicle from eight to three, measuring its accuracy under different occupancy status. Specifically, we use data collected from different combinations of UWB sensors, as described in Section 6.1, to simulate varying numbers of UWB sensors. The sensor combinations (1,2,3,4,5,6,7), (1,2,3,5,6,7), (1,3,4,5,7), (1,3,5,7), and (1,3,6) correspond to setups with seven to three UWB sensors, respectively. As shown in Figure 24, when eight sensors are used, UMusic achieves the average accuracy rates of 99.6%, 95.8%, 90.4%, and 85.8% for different occupancy status. When the number of sensors is reduced to 4, the accuracy slightly decreases but remains high at 97.2%, 93.3%, 87.5%, and 83.1%, with the decline within 3%. However, when the number of sensors is further reduced to 3, the detection performance shows a more noticeable decline, due to the reduced spatial diversity in PDP data.

6.8 Performance on Unseen Passenger
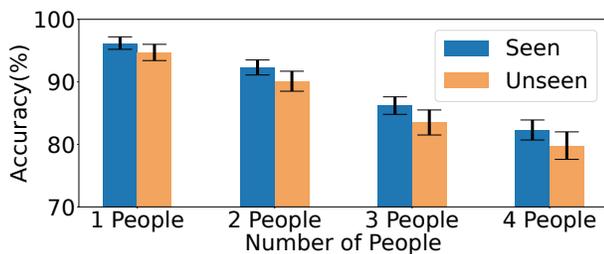
Figure 25: Performance on seen/unseen passengers.

To evaluate UMusic's performance on seen/unseen passengers, we adopt the following dataset splitting strategy: data from five randomly selected volunteers is designated as the seen dataset, while data from the remaining four volunteers constitutes the unseen dataset. The model is trained on the seen dataset, which is further split into a training set and a testing set in an 8:2 ratio. To evaluate unseen passengers, the trained model is applied to the unseen dataset, which contains data from four volunteers. Since the seen dataset includes five passengers and the unseen dataset includes four passengers, the evaluation encompasses all occupancy statuses (from 1 person to 4 people). This process is repeated for all $\binom{9}{5}$ possible combinations of seen and unseen datasets, along with the corresponding models and results. Figure 25 shows UMusic's accuracy across various in-vehicle occupancy status, where the average accuracy for seen passengers is 96.2%, 92.3%, 86.2%, and 83.5%, while the average accuracy for unseen passengers is 94.7%,
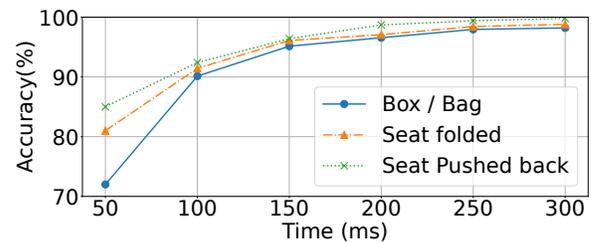
Figure 26: Performance under environment augmentation.

90.1%, 83.1%, and 78.2%, respectively. These results demonstrate that UMusic achieves high recognition accuracy even for unseen passengers, with only a slight reduction in accuracy compared to the performance on seen passengers.

6.9 Impact of Environment Augmentation

This evaluation considers three scenarios: placing a small box or bag on each seat, folding down the backseats, and pushing the front passenger seat back. To evaluate UMusic in these scenarios, we leverage the phenomenon that environmental augmentations remain static, while passengers typically exhibit unpredictable motion. By concatenating multiple consecutive CIR measurements into the SVM model, UMusic identifies occupancy status by capturing path changes caused by passengers, as described in [21]. Figure 26 illustrates UMusic's performance under environmental augmentation, with each line representing the average accuracy across four different occupancy statuses (1 to 4 people). Specifically, under the seat-pushed-back scenario, the average accuracy of UMusic is 85%, 92.4%, 96.4%, 98.7%, 99.4%, and 99.8% when one to six consecutive CIR measurements are concatenated as features for the SVM model, respectively. Meanwhile, the other two scenarios are similar.

In addition to the three scenarios that modify the in-car environment, UMusic is also capable of handling changes in the out-car environment, such as the use of sunshades. These are typically installed when passengers exit the vehicle, at which point the car is turned off. Since UMusic can detect a pet or passenger left behind within a few seconds after the engine is switched off, it does not need to operate continuously afterward. As a result, the presence of foil-like sunshades does not affect the functionality of UMusic.

6.10 PDP Calculation Precision
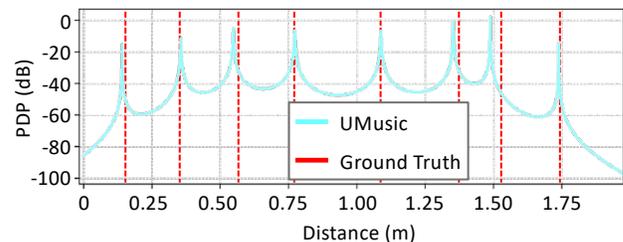
Figure 27: Comparison of UMusic and the groundtruth in the PDP, derived from the eight simulated paths.

While our evaluations demonstrate UMusic's performance in real-world scenarios, this section presents simulations to showcase PDP calculation performance in a controlled setting, where obtaining ground truth path lengths is challenging in real-world experiments. Specifically, our simulations are conducted to answer

three major questions: (i), how many paths could be precisely resolved by the high-resolution PDP calculation? Since the in-car signal propagation is very complex, manufacturers are concerning the limit of our PDP calculation. (ii), are the three hardware imperfections fully immune? The quality of UWB radio varies for different manufacturers. We should confirm the PDP is robust to these issues even for the hardware with the worst quality. (iii), how could the computational cost reduction affect the precision of PDP calculation? The onboard computer is less powerful than our laptops. It also runs many interior systems, leaving limited computational resources for UMusic. Our simulation should demonstrate the effectiveness of computational cost while maintaining the PDP calculation precision. To fit into the car's interior area (2m× 2m), our simulation focuses on paths that are less than 1.8 meters. The detailed simulation for the above three aspects is shown in the following sections.

6.10.1 Performance vs Number of Paths. In this simulation, we uniformly insert different paths into the 1.8 m compact in-car area. For instance, when we insert eight paths into the 1.8 m range, we control the average interval between two consecutive paths to be $1.8/8 = 0.225$ meters. We also consider the regulation of the FCC about the maximum SNR of UWB signal (≤ -40dB/MHz [48]) in our simulation. Figure 27 depicts an example PDP calculation result of 8 paths and the corresponding groundtruth PDP used for that simulation. This result intuitively shows that the calculation error is quite small, indicating the high precision of our design.

(a) PDP calculation error over the number of path within 1.8 m.

(b) A detailed CDF of the PDP calculation error for 7, 8 and 9 paths.
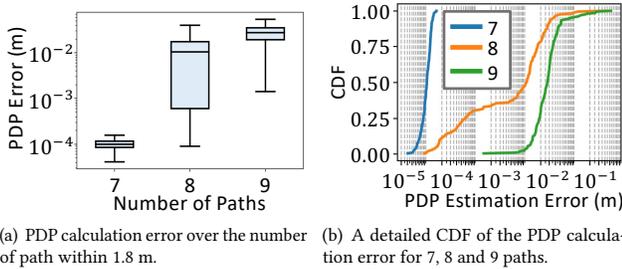
Figure 28: High-resolution PDP calculation performance vs the number of paths in small range (1.8 m).

We also summarize the simulation results into two statistics shown in Figure 28. The result in Figure 28(a) shows the PDP calculation error when we insert seven, eight, and nine paths into the 1.8 m range. It is worth mentioning that for a 500MHz bandwidth UWB channel, the 1.8 m range corresponds to only six CIR taps, which is saturated by our 7-9 paths for pushing the PDP calculation into the limit. For each result, we repeatedly generate the uniformly distributed paths 1000 times to guarantee the effectiveness of this simulation. In specific, the average interval between two consecutive paths for the seven, eight, and nine paths scenarios are $1.8/7 = 0.257$, $1.8/8 = 0.225$, and $1.8/9 = 0.2$ meters.

Due to the large gap between the two paths, the average calculation error for seven paths is 0.0001 m. When we insert eight paths, the error increase to 0.015m, while the error becomes 0.036 m if we insert nine paths. This is also validated by the CDF shown in Figure 28(b). This result shows that UMusic is able to recognize 7-9 paths which mainly reside in the six CIR taps (1.8 m).

(a) The CDF of PDP error of 8 paths under hardware imperfections.

(b) The PDP error of 8 paths under the sampling offset.

(c) The PDP error of 8 paths under CFO.

(d) The PDP error of 8 paths under the random initial phase.
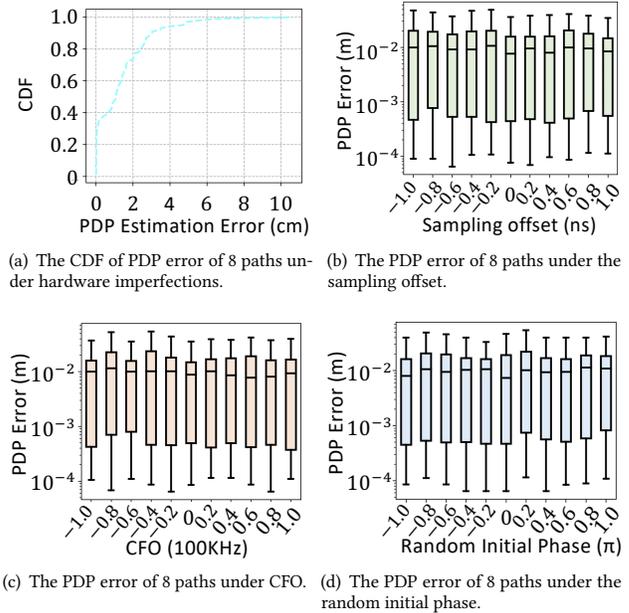
Figure 29: PDP calculation performance under three types of hardware imperfections.

6.10.2 Performance under Hardware Imperfection. Our simulation for hardware imperfection strictly follows the parameters, specifically CFO, provided by the commodity device's datasheet. As specified in DW1000 datasheet [48], the oscillator has a maximum drift of ±20ppm, resulting in the CFO maximum at the ±69.8KHz. In our simulation, we manually insert the 11 different CFO from −100KHz to 100KHz to cover all potential cases even much worse than the datasheet specifies. Moreover, we impose the sampling offset and random initial phase offset to the simulated CIR data to complete the verification for hardware immunity. CFO, initial offset, random sampling time offset. The overall PDP calculation error is shown in Figure 29(a), where the average error is 0.016 m, which matches with the 8 paths calculation result in Section 6.10.1 obtained under no hardware imperfections. In specific, we show the impact of each hardware imperfection in Figure 29(b), 29(c), and 29(d), respectively. These results demonstrate that the distribution of PDP calculation error is constant even with various hardware imperfections, which verifies the immunity of our high-resolution PDP calculation.

(a) Time consumed for PDP calculation w.r.t. the downsampling factor D.

(b) The CDF of PDP estimation error w.r.t. the the downsampling factor D.
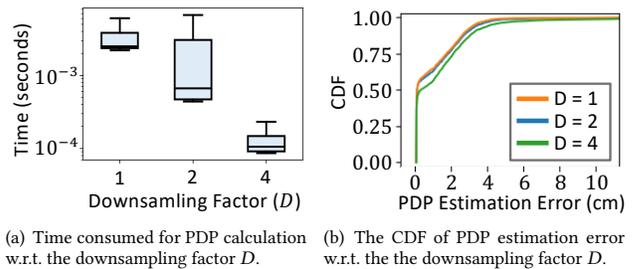
Figure 30: Results of computational cost optimization.

6.10.3 With/Without Computational Cost Reduction. We evaluate the time consumption for high-resolution PDP calculation. In this

experiment, we reuse the simulated CIR data and their corresponding groundtruth PDP information for eight paths within 1.8 m range. We implement the downsampling design to reduce the computational cost on a Mac M1 computer, which is set to use a single CPU and a single thread for PDP calculation. This experiment compares the PDP calculation time and precision under the downsampling factors 1, 2, and 4. When we downsample by a factor of 4, we set p and q to be 48 in order to make them the multipliers of $D = 4$. The detailed results are shown in Figure 30, where the average time consumption is 3.38 ms, 1.15 ms, and 0.125 ms under the downsampling factor of 1, 2, and 4 respectively. Compared to $D = 1$, the time consumption is reduced by 2.94 and 27.1 times. Meanwhile, the accuracy experiences a negligible degradation due to downsampling, as shown in Figure 30(b). This result supports UMusic to operate in real-time since the interval between consecutive CIR collections is 50 ms. As UMusic is intended for in-car systems powered by the alternator, its energy consumption of UMusic is relatively low compared to other in-car systems.

7 DISCUSSION AND FUTURE WORK

Impact of tall passenger. In the experiment, UWB devices are mounted on the car's ceiling to minimize the impact of tall passengers. Even if an exceptionally tall passenger blocks the LoS path, UMusic can handle this effectively. Although PDP synchronization may be affected, this does not hinder occupancy detection, as the blockage also impacts longer paths, making the shortest affected path closer in length to the LoS path. Subsequently, this could be captured by UMusic to detect the exceptionally tall passenger.

Distinguishability between the passenger and large luggage. We note that luggage remains stationary, while passengers typically display unpredictable motion. This distinction allows UMusic to differentiate large luggage from passengers by tracking changes in propagation paths over time, reusing PDP data collected for aggregated detection, as studied in [21, 32, 87].

Extension to general sensing applications. While UMusic is tailored for in-car occupancy detection due to strong signal reflections from the vehicle's metal and compact structure, it can also be adapted for broader sensing applications. Its core feature, high-resolution PDP estimation, can be utilized in tasks like localization. By using multiple UWB sensors as anchors, the precise PDP of sensor links can enable accurate triangulation for localization.

Full support for HVAC and vital sign applications. UMusic is designed to provide occupancy status, a prerequisite for HVAC systems, vital sign monitoring, and detect children left alone. Currently, UMusic excels at occupancy detection, identifying the number of passengers and their seating arrangements, which is suitable for HVAC applications. However, UMusic requires further enhancements to detect physiological signals such as heart rate, respiration rate, and body temperature for vital sign monitoring [83, 91] and other applications, which would be addressed in future work.

8 RELATED WORK

Wireless sensing has been studied in numerous papers [27, 73] for indoor localization [24, 56, 72, 80], location tracking [47, 73, 79], floor mapping [31, 46], and motion tracking [84] . These designs, which are categorized by the wireless techniques utilized, have different advantages and limitations: (i), WiFi-based designs [24, 47, 56, 66, 68, 78] are easy for users to accept as the result of the popularity of WiFi [69, 70]. (ii), Acoustic [7, 90] and vision [12] based approaches have high precision while the privacy concerns are yet to be addressed. (iii), mmWave radar is also applied for sensing [30, 31, 36, 42, 58, 67, 73] and achieved both strong privacy reservation and high effectiveness. (iv), UWB is recently applied for precise localization [3, 43, 57, 60, 61, 63] and sensing [21, 32]. For instance, TALLA [63] achieves decimal-level localization and tracking precision using time difference of arrival (TDoA), derived from UWB communication. (v), Bluetooth [33, 65], LoRa [6, 15, 27, 74, 77], RFID [10, 75, 76] are also leveraged for sensing in various scenarios, where the granularity is not strictly required.

For in-car scenarios, mmWave Radar [28, 30, 31, 42, 46, 73], vision-based [12] and acoustic-based [90] approaches have high sensitivity for occupancy detection. Vision-based solutions suffer from occlusions, while acoustic-based solutions still face privacy leakage issues. For instance, VeCare proposes the first Child Presence Detection (CPD) system that only utilizes car speakers and microphones. To ensure a robust solution, these approaches require additional hardware and associated installation costs. In addition, customizing a low-cost tag is an effective solution for in-car sensing [8]. This paper leverages UWB devices installed in the car for occupancy detection. Existing UWB sensing solutions [21, 32] have primarily focused on utilizing machine learning techniques, making them sensitive to changes in the in-car environment. For instance, CarOSense investigates the reuse of UWB keyless infrastructure through a novel deep-learning model called MaskMIMO to detect occupancy in each seat of a car [32]. UMusic, on the other hand, utilizes signal processing techniques making the solution more adaptive to the environmental effects. A combination of advanced ML techniques with UMusic can potentially create more robust hybrid models to detect and classify multiple occupancies in the various car models. Development and analysis of such models are left for future work.

9 CONCLUSION

This paper introduces UMusic, a system that uses commodity UWB devices to precisely detect car occupancy via lightweight signal processing techniques. UMusic converts CIR data into the frequency domain to obtain the channel frequency response, which is used to calculate the high-resolution PDP via the MUSIC algorithm. Through the comparison between the PDP of empty and occupied environments, UMusic is able to detect the occupancy status. We evaluate UMusic in a car with one or more passengers under various scenarios, including stationary and driving conditions. The experiments show that UMusic achieves an aggregated accuracy of 99.4%, highlighting its effectiveness in practical scenarios.

REFERENCES

[1] 4activeSystems. In-Cabin Dummies, https://www.4activesystems.at/, 2024.

[2] BMW. Introducing the BMW Digital Key Plus, https://www.bmw.com/en/innovation/bmw-digital-key-plus-ultra-wideband.html, 2021.

[3] Y. Cao, A. Dhekne, and M. Ammar. Itracku: tracking a pen-like instrument via uwb-imu fusion. In *Proceedings of the 19th Annual International Conference on Mobile Systems, Applications, and Services*, pages 453–466, 2021.

[4] M. Cellario. Human-centered intelligent vehicles: Toward multimodal interface integration. *IEEE intelligent systems*, 16(4):78–81, 2001.

[5] L. Chang-won. Hyundai auto group provides new digital key service for Samsung and Apple smartphones, https://www.ajudaily.com/view/20220412113249182, 2022.

[6] W. Chen, J. Zhang, X. Xia, S. Wang, and T. He. Hitting the sweet spot: An sf-any coding paradigm for empowering city-wide lora communications. In *2024 23rd ACM/IEEE International Conference on Information Processing in Sensor Networks (IPSN)*, pages 225–236. IEEE, 2024.

[7] Y. Chen, W. Gong, J. Liu, and Y. Cui. I can hear more: Pushing the limit of ultrasound sensing on off-the-shelf mobile devices. In *IEEE INFOCOM 2018-IEEE Conference on Computer Communications*, pages 2015–2023. IEEE, 2018.

[8] M. Cui, B. Xie, Q. Wang, and J. Xiong. Evlesen: In-vehicle sensing with ev-leaked signal. In *Proceedings of the 30th Annual International Conference on Mobile Computing and Networking*, ACM MobiCom '24, page 679–693, New York, NY, USA, 2024. Association for Computing Machinery.

[9] M. De Sanctis, A. Conte, T. Rossi, S. Di Domenico, and E. Cianca. Cir-based device-free people counting via uwb signals. *Sensors*, 21(9):3296, 2021.

[10] H. Ding, L. Shangguan, Z. Yang, J. Han, Z. Zhou, P. Yang, W. Xi, and J. Zhao. Femo: A platform for free-weight exercise monitoring with rfids. In *Proceedings of the 13th ACM conference on embedded networked sensor systems*, pages 141–154, 2015.

[11] S. Electronics. Remote Keyless Entry Schemes for an Automobile Using the DecaWave DW1000, https://www.symmetryelectronics.com/blog/remote-keyless-entry-schemes-for-an-automobile-using-the-decawave-dw1000/?srsltid=afmbooopr1j6fd365ublz9luvstndqcnsljhotipvup6oofau9ojygze, 2014.

[12] F. Erlik Nowruzi, W. A. El Ahmar, R. Laganiere, and A. H. Ghods. In-vehicle occupancy detection with convolutional networks on thermal images. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops*, pages 0–0, 2019.

[13] Ford. What is Ford Pre-Collision Assist with Pedestrian Detection ?, https://www.ford.com/support/how-tos/ford-technology/driver-assist-features/what-is-pre-collision-assist-with-pedestrian-detection/.

[14] L. Fridman. Human-centered autonomous vehicle systems: Principles of effective shared autonomy. *arXiv preprint arXiv:1810.01835*, 2018.

[15] D. Gao, H. Wang, S. Wang, W. Wang, Z. Yin, S. Mumtaz, X. Li, V. Frascolla, and A. Nallanathan. Wilo: Long-range cross-technology communication from wi-fi to lora. *IEEE Transactions on Communications*, 2024.

[16] A. Gharamohammadi, A. Khajepour, and G. Shaker. In-vehicle monitoring by radar: A review. *IEEE Sensors Journal*, 2023.

[17] K. Golsch and S. Sute. Passive entry/passive start systems and methods for vehicles, June 25 2019. US Patent 10,328,898.

[18] L. Hogben. *Handbook of linear algebra*. CRC press, 2013.

[19] Honda. Honda CR-V Earns Top Pedestrian Detection and Crash Prevention Ratings, https://csr.honda.com/2019/03/01/honda-cr-v-earns-top-pedestrian-detection-and-crash-prevention-ratings/, 2019.

[20] S. S. Jayaweera, B. Wang, X. Zeng, W.-H. Wang, and K. R. Liu. Wifi-based robust child presence detection for smart cars. In *ICASSP 2023-2023 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 1–5. IEEE, 2023.

[21] A. Kalyanaraman, Y. Zeng, S. Rakshit, and V. Jain. Caraokey: Car states sensing via the ultra-wideband keyless infrastructure. In *2020 17th Annual IEEE International Conference on Sensing, Communication, and Networking (SECON)*, pages 1–9. IEEE, 2020.

[22] J. S. Katrolia, B. Mirbach, A. El-Sherif, H. Feld, J. Rambach, and D. Stricker. Ticam: A time-of-flight in-car cabin monitoring dataset. *arXiv preprint arXiv:2103.11719*, 2021.

[23] K. S. Kenney and L. M. Fanciullo. Automobile air bags: friend or foe? a case of air bag-associated ocular trauma and a related literature review. *Optometry-Journal of the American Optometric Association*, 76(7):382–386, 2005.

[24] M. Kotaru, K. Joshi, D. Bharadia, and S. Katti. Spotfi: Decimeter level localization using wifi. In *Proceedings of the 2015 ACM Conference on Special Interest Group on Data Communication*, pages 269–282, 2015.

[25] A. Lazaro, M. Lazaro, R. Villarino, and D. Girbau. Seat-occupancy detection system and breathing rate monitoring based on a low-cost mm-wave radar at 60 ghz. *IEEE Access*, 9:115403–115414, 2021.

[26] T. Lei, S. Pan, Q. Guang, K. Wang, and Y. Yu. Enhanced geometric filtering method based device-free localization with uwb wireless network. *IEEE Transactions on Vehicular Technology*, 70(8):7734–7748, 2021.

[27] J. Liu, J. Gao, S. Jha, and W. Hu. Seirios: leveraging multiple channels for lorawan indoor and outdoor localization. In *Proceedings of the 27th Annual International Conference on Mobile Computing and Networking*, pages 656–669, 2021.

[28] R. Liu, T. Yao, R. Shi, L. Mei, S. Wang, Z. Yin, W. Jiang, and S. Wang. Mission: mmwave radar person identification with rgb cameras. In *Proceedings of the 22nd ACM Conference on Embedded Networked Sensor Systems*, pages 309–321, 2024.

[29] S. Liu, K. Koch, Z. Zhou, S. Föll, X. He, T. Menke, E. Fleisch, and F. Wortmann. The empathetic car: Exploring emotion inference via driver behaviour and traffic context. *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies*, 5(3):1–34, 2021.

[30] C. X. Lu, S. Rosa, P. Zhao, B. Wang, C. Chen, J. A. Stankovic, N. Trigoni, and A. Markham. See through smoke: robust indoor mapping with low-cost mmwave radar. In *Proceedings of the 18th International Conference on Mobile Systems, Applications, and Services*, pages 14–27, 2020.

[31] C. X. Lu, M. R. U. Saputra, P. Zhao, Y. Almalioglu, P. P. De Gusmao, C. Chen, K. Sun, N. Trigoni, and A. Markham. milliego: single-chip mmwave radar aided egomotion estimation via deep sensor fusion. In *Proceedings of the 18th Conference on Embedded Networked Sensor Systems*, pages 109–122, 2020.

[32] Y. Ma, Y. Zeng, and V. Jain. Carosense: Car occupancy sensing with the ultra-wideband keyless infrastructure. *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies*, 4(3):1–28, 2020.

[33] E. Mackensen, M. Lai, and T. M. Wendt. Bluetooth low energy (ble) based wireless sensors. In *SENSORS, 2012 IEEE*, pages 1–4. IEEE, 2012.

[34] G. Marrow, M. McLaughlin, and C. McElroy. Measuring angle of incidence in an ultrawideband communication system, Sept. 11 2018. US Patent 10,075,210.

[35] Mazda. Driver and Front Passenger Occupant Classification System, https://www.mazdausa.com/static/manuals/2023/cx-30/contents/03040702.html#:~:text=your%20vehicle%20is%20equipped%20withthe%20front%20passenger%20seat%20rails., 2023.

[36] L. Mei, S. Wang, Y. Cheng, R. Liu, Z. Yin, W. Jiang, and W. Gong. Esp-pct: Enhanced vr semantic performance through efficient compression of temporal and spatial redundancies in point cloud transformers. In *33rd International Joint Conference on Artificial Intelligence (IJCAI 2024)*, pages 1182–1190. International Joint Conferences on Artificial Intelligence, 2024.

[37] A. Mishra, S. Lee, D. Kim, and S. Kim. In-cabin monitoring system for autonomous vehicles. *Sensors*, 22(12), 2022.

[38] E. NCAP. Safe Driving Occupant Monitoring, https://www.euroncap.com/media/83735/euro-ncap-protocol-safe-driving-occupant-monitoring-v09.pdf, 2024.

[39] N. H. T. S. A. (NHTSA). Federal Motor Vehicle Safety Standards; Occupant Crash Protection, Seat Belt Reminder Systems, https://www.federalregister.gov/documents/2023/09/07/2023-18413/federal-motor-vehicle-safety-standards-occupant-crash-protection-seat-belt-reminder-systems, 2023.

[40] NXP. NXP and VW share the wide possibilities of Ultra-Wideband's (UWB) fine ranging capabilities, https://www.nxp.com/company/about-nxp/nxp-and-vw-share-the-wide-possibilities-of-ultra-widebands-uwb-fine-ranging-capabilities:nw-volkswagen-showcases-uwb, 2019.

[41] A. V. Oppenheim. *Discrete-time signal processing*. Pearson Education India, 1999.

[42] M. Z. Ozturk, C. Wu, B. Wang, and K. R. Liu. Toward mmwave-based sound enhancement and separation. In *ICASSP 2022-2022 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 6852–6856. IEEE, 2022.

[43] P. Pannuto, B. Kempke, and P. Dutta. Slocalization: Sub-uw ultra wideband backscatter localization. In *2018 17th ACM/IEEE International Conference on Information Processing in Sensor Networks (IPSN)*, pages 242–253. IEEE, 2018.

[44] F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, et al. Scikit-learn: Machine learning in python. *Journal of machine learning research*, 12(Oct):2825–2830, 2011.

[45] T. Peto. PyArgus: a python package aims to implement signal processing algorithms applicable in antenna arrays., https://github.com/petotamas/pyargus, 2021.

[46] A. Prabhakara, T. Jin, A. Das, G. Bhatt, L. Kumari, E. Soltanaghaei, J. Bilmes, S. Kumar, and A. Rowe. High resolution point clouds from mmwave radar. *arXiv preprint arXiv:2206.09273*, 2022.

[47] K. Qian, C. Wu, Y. Zhang, G. Zhang, Z. Yang, and Y. Liu. Widar2. 0: Passive human tracking with a single wi-fi link. In *Proceedings of the 16th Annual International Conference on Mobile Systems, Applications, and Services*, pages 350–361, 2018.

[48] Qorvo. DW1000 Datasheet, https://www.qorvo.com/products/d/da007946, 2017.

[49] Qorvo. DW3000 Datasheet, https://www.qorvo.com/products/d/da008142, 2020.

[50] radino32DW1000. radino32 DW1000 for Ranging and RTLS, Ultra Wideband (UWB), IEEE802.15.4-2011 compliant, 3.5-6.5GHz, https://shop.in-circuit.de/product_info.php?products_id=186&language=en, 2021.

[51] M. Rinta-Homi, N. H. Motlagh, A. Zuniga, H. Flores, and P. Nurmi. How low can you go? performance trade-offs in low-resolution thermal sensors for occupancy detection: A systematic evaluation. *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies*, 5(3):1–22, 2021.

[52] R. Schmidt. Multiple emitter location and signal parameter estimation. *IEEE transactions on antennas and propagation*, 34(3):276–280, 1986.

[53] N. Semiconductors. Determine a Device's Relative Position with UWB Localization Capability, https://www.nxp.com/applications/enabling-technologies/connectivity/ultra-wideband-uwb:uwb, 2021.

[54] T.-J. Shan, M. Wax, and T. Kailath. On spatial smoothing for direction-of-arrival estimation of coherent signals. *IEEE Transactions on Acoustics, Speech, and Signal Processing*, 33(4):806–811, 1985.

[55] M. Sidikova, R. Martinek, A. Kawala-Sterniuk, M. Ladrova, R. Jaros, L. Danys, and P. Simonik. Vital sign monitoring in car seats based on electrocardiography, ballistocardiography and seismocardiography: A review. *Sensors*, 20(19):5699, 2020.

[56] E. Soltanaghaei, A. Kalyanaraman, and K. Whitehouse. Multipath triangulation: Decimeter-level wifi localization and orientation with a single unaided receiver. In *Proceedings of the 16th annual international conference on mobile systems, applications, and services*, pages 376–388, 2018.

[57] K. Sun, S. Wang, R. Zhao, R. Liu, W. Chen, Z. Yin, W. Gong, and S. Wang. Wuloc: Achieving extremely long-range high-precision localization via wi-fi-uwb connection. *Proc. ACM Interact. Mob. Wearable Ubiquitous Technol.*, 9(1), Mar. 2025.

[58] M. Tang, P. Teckchandani, J. He, H. Guo, and E. Soltanaghai. Bsense: In-vehicle child detection and vital sign monitoring with a single mmwave radar and synthetic reflectors. In *Proceedings of the 22nd ACM Conference on Embedded Networked Sensor Systems*, pages 478–492, 2024.

[59] Tesla. Tesla Model 3 Collision Avoidance Assist, https://www.tesla.com/ownersmanual/model3/en_us/guid-8ea7ef10-7d27-42ac-a31a-96bce5bc0a85.html, 2022.

[60] J. Tiemann, F. Eckermann, and C. Wietfeld. Atlas-an open-source tdoa-based ultra-wideband localization system. In *2016 International Conference on Indoor Positioning and Indoor Navigation (IPIN)*, pages 1–6. IEEE, 2016.

[61] J. Tiemann, Y. Elmasry, L. Koring, and C. Wietfeld. Atlas fast: Fast and simple scheduled tdoa for reliable ultra-wideband localization. In *2019 International Conference on Robotics and Automation (ICRA)*, pages 2554–2560. IEEE, 2019.

[62] M. Vamsi and K. Soman. In-vehicle occupancy detection and classification using machine learning. In *2020 11th International Conference on Computing, Communication and Networking Technologies (ICCCNT)*, pages 1–6. IEEE, 2020.

[63] D. Vecchia, P. Corbalán, T. Istomin, and G. P. Picco. Talla: Large-scale tdoa localization with ultra-wideband radios. In *2019 International Conference on Indoor Positioning and Indoor Navigation (IPIN)*, pages 1–8. IEEE, 2019.

[64] F. Wang, J. Liu, and W. Gong. Wicar: Wifi-based in-car activity recognition with multi-adversarial domain adaptation. In *Proceedings of the International Symposium on Quality of Service*, pages 1–10, 2019.

[65] H. Wang, J. Wang, W. Jiang, S. Wang, and D. Gao. Physical layer cross-technology communication via explainable neural networks. *IEEE Transactions on Mobile Computing*, 2024.

[66] J. Wang, H. Jiang, J. Xiong, K. Jamieson, X. Chen, D. Fang, and B. Xie. Lifs: Low human-effort, device-free localization with fine-grained subcarrier information. In *Proceedings of the 22nd Annual International Conference on Mobile Computing and Networking*, pages 243–256, 2016.

[67] S. Wang, D. Cao, R. Liu, W. Jiang, T. Yao, and C. X. Lu. Human parsing with joint learning for dynamic mmwave radar point cloud. *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies*, 7(1):1–22, 2023.

[68] S. Wang, J. Guo, P. Wang, K. Parsons, P. Orlik, Y. Nagai, T. Sumi, and P. Pathak. X-disco: Cross-technology neighbor discovery. In *2022 19th Annual IEEE International Conference on Sensing, Communication, and Networking (SECON)*, pages 163–171. IEEE, 2022.

[69] S. Wang, W. Jeong, J. Jung, and S. M. Kim. X-mimo: Cross-technology multi-user mimo. In *Proceedings of the 18th Conference on Embedded Networked Sensor Systems*, pages 218–231, 2020.

[70] S. Wang, S. M. Kim, and T. He. Symbol-level cross-technology communication via payload encoding. In *2018 IEEE 38th International Conference on Distributed Computing Systems (ICDCS)*, pages 500–510. IEEE, 2018.

[71] Z. Wang, Z. Chen, A. D. Singh, L. Garcia, J. Luo, and M. B. Srivastava. Uwhear: through-wall extraction and separation of audio vibrations using wireless signals. In *Proceedings of the 18th Conference on Embedded Networked Sensor Systems*, pages 1–14, 2020.

[72] C. Wu, Z. Yang, Y. Liu, and W. Xi. Will: Wireless indoor localization without site survey. *IEEE Transactions on Parallel and Distributed systems*, 24(4):839–848, 2012.

[73] C. Wu, F. Zhang, B. Wang, and K. J. R. Liu. mmtrack: Passive multi-person localization using commodity millimeter wave radio. In *IEEE International Conference on Computer Communications*, April 27-30 2020.

[74] B. Xie and J. Xiong. Combating interference for long range lora sensing. In *Proceedings of the 18th Conference on Embedded Networked Sensor Systems*, pages 69–81, 2020.

[75] B. Xie, J. Xiong, X. Chen, E. Chai, L. Li, Z. Tang, and D. Fang. Tagtag: material sensing with commodity rfid. In *Proceedings of the 17th conference on embedded networked sensor systems*, pages 338–350, 2019.

[76] B. Xie, J. Xiong, X. Chen, and D. Fang. Exploring commodity rfid for contactless sub-millimeter vibration sensing. In *Proceedings of the 18th Conference on Embedded Networked Sensor Systems*, pages 15–27, 2020.

[77] B. Xie, Y. Yin, and J. Xiong. Pushing the limits of long range wireless sensing with lora. *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies*, 5(3):1–21, 2021.

[78] Y. Xie, Z. Li, and M. Li. Precise power delay profiling with commodity wifi. In *Proceedings of the 21st Annual International Conference on Mobile Computing and Networking*, pages 53–64, 2015.

[79] Y. Xie, J. Xiong, M. Li, and K. Jamieson. md-track: Leveraging multi-dimensionality for passive indoor wi-fi tracking. In *The 25th Annual International Conference on Mobile Computing and Networking*, pages 1–16, 2019.

[80] J. Xiong and K. Jamieson. Arraytrack: A fine-grained indoor location system. Usenix, 2013.

[81] Q. Xu, B. Wang, F. Zhang, D. S. Regani, F. Wang, and K. R. Liu. Wireless ai in smart car: How smart a car can be? *IEEE Access*, 8:55091–55112, 2020.

[82] J. Yang, M. Wang, Y. Yang, and J. Wang. A device-free localization and size prediction system for road vehicle surveillance via uwb networks. *IEEE Transactions on Instrumentation and Measurement*, 71:1–11, 2021.

[83] Z. Yang, M. Bocca, V. Jain, and P. Mohapatra. Contactless breathing rate monitoring in vehicle using uwb radar. In *Proceedings of the 7th international workshop on real-world embedded wireless systems and networks*, pages 13–18, 2018.

[84] S. Yun, Y.-C. Chen, H. Zheng, L. Qiu, and W. Mao. Strata: Fine-grained acoustic-based device-free tracking. In *Proceedings of the 15th annual international conference on mobile systems, applications, and services*, pages 15–28, 2017.

[85] X. Zeng, B. Wang, C. Wu, S. D. Regani, and K. R. Liu. Wicpd: Wireless child presence detection system for smart cars. *IEEE Internet of Things Journal*, 9(24):24866–24881, 2022.

[86] X. Zeng, F. Wang, B. Wang, C. Wu, K. R. Liu, and O. C. Au. In-vehicle sensing for smart cars. *IEEE Open Journal of Vehicular Technology*, 2022.

[87] Y. Zeng and V. Jain. Methods and systems for estimating vital signs of vehicle occupants using ultra wideband, Nov. 12 2024. US Patent 12,139,087.

[88] F. Zhang, J. Xiong, Z. Chang, J. Ma, and D. Zhang. Mobi2sense: empowering wireless sensing with mobility. In *Proceedings of the 28th Annual International Conference on Mobile Computing And Networking*, pages 268–282, 2022.

[89] S. Zhang, T. Zheng, Z. Chen, and J. Luo. Can we obtain fine-grained heartbeat waveform via contact-free rf-sensing? In *IEEE INFOCOM 2022-IEEE Conference on Computer Communications*, pages 1759–1768. IEEE, 2022.

[90] Y. Zhang, W. Hou, Z. Yang, and C. Wu. {VeCare}: Statistical acoustic sensing for automotive {In-Cabin} monitoring. In *20th USENIX Symposium on Networked Systems Design and Implementation (NSDI 23)*, pages 1185–1200, 2023.

[91] T. Zheng, Z. Chen, C. Cai, J. Luo, and X. Zhang. V2ifi: In-vehicle vital sign monitoring via compact rf sensing. *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies*, 4(2):1–27, 2020.

[92] T. Zheng, Z. Chen, S. Zhang, C. Cai, and J. Luo. More-fi: Motion-robust and fine-grained respiration monitoring via deep-learning uwb radar. In *Proceedings of the 19th ACM Conference on Embedded Networked Sensor Systems*, pages 111–124, 2021.