SpecRing: Spectrum-Shifting Encryption for Privacy-Preserving mmWave Sensing

HANKAI LIU, Nankai University, China
JIANYANG WANG✉, University of Science and Technology of China, China
RUI ZHANG, University of Science and Technology of China, China
RUIXU GENG, The University of Hong Kong, China
MENG XI, Tianjin University, China
XIULONG LIU, Tianjin University, China
DAN DING, Nankai University, China
YAN CHEN, University of Science and Technology of China, China

Millimeter-wave (mmWave) sensing supports a broad range of contactless applications for ubiquitous computing. Despite being regarded as a privacy-friendly modality, raw radar streams encode sensitive physical, behavioral, and physiological semantics that can be exploited by unauthorized analytics. Existing protections either force an all-or-nothing trade-off between confidentiality and utility, or apply task-specific perturbations that provide neither cryptographic access control nor cross-task generality. To bridge this gap, this paper presents SpecRing, a spectrum-shifting encryption framework that protects mmWave data at the semantic level while preserving radar-processable signal structure. SpecRing exploits the separable dimensions of mmWave signals, and applies phase offsets to cyclically rotate selected spectra. Authorized services can apply the inverse shifts to recover the true physical scene, whereas unauthorized parties observe a format-valid but semantically misleading virtual scene. To make this design practical, SpecRing removes environmental imprints that could reveal secret offsets, and generates server-specific offsets through pairwise key negotiation and deterministic derivation. Experiments show that SpecRing achieves exact recovery with correct keys, produces random-guessing-level errors with wrong keys, reduces environmental-imprint attack success rates to below 1%, and incurs 25.6 ms user-side and 10.3 ms server-side overhead per frame. Case studies on tracking, classification, vital-sign monitoring, and radar imaging demonstrate its broad applicability.

CCS Concepts: • **Human-centered computing** → **Ubiquitous and mobile computing**; • **Security and privacy** → **Security services**.

Additional Key Words and Phrases: mmWave Sensing, Privacy-preserving Sensing, Wireless Data Security

ACM Reference Format:
Hankai Liu, Jianyang Wang, Rui Zhang, Ruixu Geng, Meng Xi, Xiulong Liu, Dan Ding, and Yan Chen. 2026. SpecRing: Spectrum-Shifting Encryption for Privacy-Preserving mmWave Sensing. *Proc. ACM Interact. Mob. Wearable Ubiquitous Technol.* 10, 4, Article 240 (December 2026), 28 pages. https://doi.org/10.1145/3857968

Authors' Contact Information: Hankai Liu, Nankai University, Tianjin, China, lhk@nankai.edu.cn; Jianyang Wang (corresponding author), University of Science and Technology of China, Hefei, China, wangjianyang@mail.ustc.edu.cn; Rui Zhang, University of Science and Technology of China, Hefei, China, ruizh@mail.ustc.edu.cn; Ruixu Geng, The University of Hong Kong, Hong Kong, China, raygeng@hku.hk; Meng Xi, Tianjin University, Tianjin, China, ximeng@tju.edu.cn; Xiulong Liu, Tianjin University, Tianjin, China, xiulong_liu@tju.edu.cn; Dan Ding, Nankai University, Tianjin, China, dingdan@nankai.edu.cn; Yan Chen, University of Science and Technology of China, Hefei, China, eecyan@ustc.edu.cn.



This work is licensed under a Creative Commons Attribution-NonCommercial-NoDerivatives 4.0 International License.
© 2026 Copyright held by the owner/author(s).
ACM 2474-9567/2026/12-ART240
https://doi.org/10.1145/3857968

1 Introduction

1.1 Motivation and Problem Statement

Millimeter-wave (mmWave) sensing is becoming an increasingly attractive modality for ubiquitous computing. By analyzing radio reflections from the human body and surrounding environments, commodity mmWave radars support a broad range of contactless applications, including human tracking [20], activity recognition [55], gesture interaction [22], fall detection [28], and vital signs monitoring [36, 48]. Compared with cameras and microphones, mmWave sensors do not directly capture images or audio, and thus are often regarded as a privacy-friendly alternative for in-home and healthcare sensing.

This perception, however, is only partially true [41]. The same signal properties that make mmWave useful for human sensing also make raw radar streams privacy-sensitive [12]. In other words, mmWave sensing does not avoid privacy exposure. Instead, it shifts that exposure from directly observable content to latent physical semantics. This creates a subtle but important privacy risk. Precisely because mmWave sensing appears less intrusive, users may be more willing to deploy such devices and share their data [10]. A practical source of this risk is outsourced sensing. A radar owner may provide radar data to a third-party service for a specific task, such as localization or health monitoring. Although the service is authorized to process the stream for that task, it should not infer unrelated physical, behavioral, or physiological information [16]. However, access to the raw stream makes such secondary inference technically possible. Such access can reveal in-home routines, behavioral biometrics, health-related conditions, and speech content [19, 39, 40, 43]. The problem, therefore, is not whether mmWave sensing is useful, but how to preserve its sensing utility while preventing the underlying physical semantics from being exposed to unauthorized parties.

1.2 Limitations of Prior Art

Existing privacy mechanisms do not fully address this problem. Conventional encryption protects radar data during transmission or storage, but cannot restrict how an authorized service uses the data after decryption. Once the stream is decrypted for the intended sensing task, the service can reuse the same data to infer information beyond its authorization. This all-or-nothing model is poorly suited to ubiquitous sensing deployments, where different services may require different physical information from the same sensor stream. For example, a localization service may need distance and angle, while a health service may rely on subtle phase variations for vital signs estimation.

Prior mmWave sensing defenses modify physical signals or sensing data. Physical-layer methods use crafted signals [15] or engineered reflectors [29, 34] to alter radar observations, but may also affect authorized sensing services. At the data level, privacy-preserving transformations [24, 33] and federated learning [31, 53] reduce privacy leakage, but remain tied to predefined attributes, tasks, or training pipelines. MuFilter [23] exploits FFT invertibility to selectively perturb predefined behavioral signatures in multi-user scenarios. However, its unencrypted output allows a communication eavesdropper to infer any information not explicitly removed. The method therefore provides neither server-specific authorization nor cryptographic confidentiality.

1.3 SpecRing in a Nutshell

This paper asks: Can we protect mmWave sensing data at the level of physical semantics, while preserving the signal structure needed by authorized sensing tasks?

We answer this question with SpecRing, a spectrum-shifting encryption framework for privacy-preserving mmWave sensing. SpecRing is based on a key observation: the raw frequency-modulated continuous-wave (FMCW) radar signal is naturally organized as a three-dimensional complex matrix, where analog-to-digital converter (ADC) samples, chirps, and antennas correspond to range, Doppler, and angle spectra, respectively. These dimensions are physically meaningful and mathematically separable. Therefore, instead of encrypting the

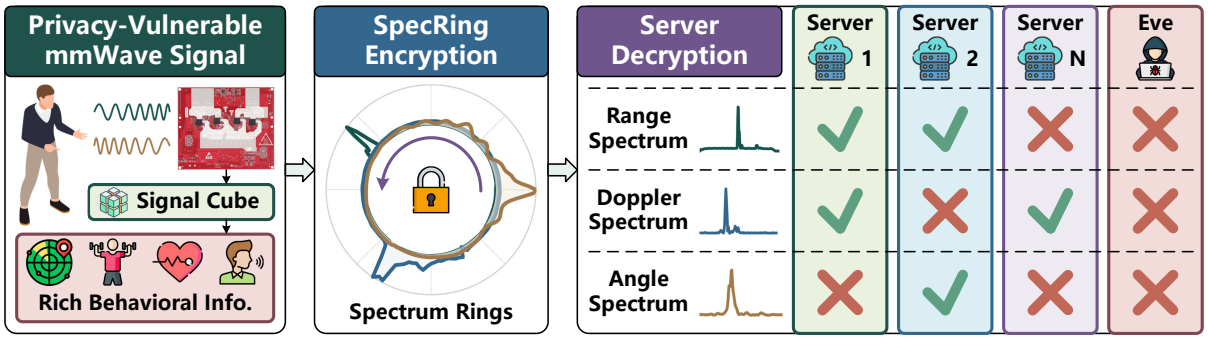


Fig. 1. Deployment concept of SpecRing.

entire signal as an opaque byte stream, SpecRing applies secret phase offsets along selected signal dimensions. In the frequency domain, this operation rotates the corresponding spectrum on a cyclic ring, as shown in Fig. 1. An authorized party holding the correct key can apply the inverse rotation and recover the original range, velocity, or angle. An unauthorized party, however, observes a format-valid radar stream whose spectral peaks are shifted to arbitrary locations. This property is particularly useful for ubiquitous sensing, where privacy violations often arise from opportunistic reuse of sensor streams by untrusted services rather than from explicit cryptographic decryption attempts.

1.4 Design Challenges

A capable adversary may possess sufficient knowledge of the proposed algorithms, RF sensing expertise, and user-server communication eavesdropping capabilities. To defend against such threats, the implementation of SpecRing must address the following challenges.

Challenge 1: How to encrypt signal spectra without disrupting their use in authorized sensing services? A signal dimension not essential to one task may still support others. For example, multi-person vital signs monitoring relies on range and angle to separate different individuals, while occupancy detection may use velocity as auxiliary evidence for moving targets. Thus, simply tampering with or masking a dimension can degrade authorized sensing utility. SpecRing addresses this challenge through spectrum-shifting encryption, which changes the failure mode of unauthorized sensing. It preserves the signal structure of the encrypted dimension while changing its physical semantics using secret offsets. Unauthorized parties can still conduct signal processing, but their results correspond to a secret-induced virtual scene rather than the real environment. Authorized services, with the proper key, can recover the true semantics and operate normally.

Challenge 2: How to prevent environmental imprints from revealing secret offsets? A direct implementation of spectrum shifting is vulnerable because it shifts spectra not only of the target but also of stable background structures. Indoor radar signals contain environmental imprints such as static reflector peaks from walls and furniture, boundary-bin energy valleys caused by transmit leakage and static filtering, and attenuation profiles induced by path loss. Since these imprints are shifted by the same secret offsets as the target signals, a knowledgeable adversary may use them as anchors to infer the secret offsets. SpecRing addresses this challenge by removing such anchors from the encrypted spectrum. It suppresses the signal energy of static structures and then performs frequency-domain processing using unoccupied-room reference signals, including background range bin substitution and zero-velocity Doppler bin superposition. As a result, the encrypted spectrum no longer contains stable environmental imprints, preventing adversaries from estimating the secret offsets from background spectra.

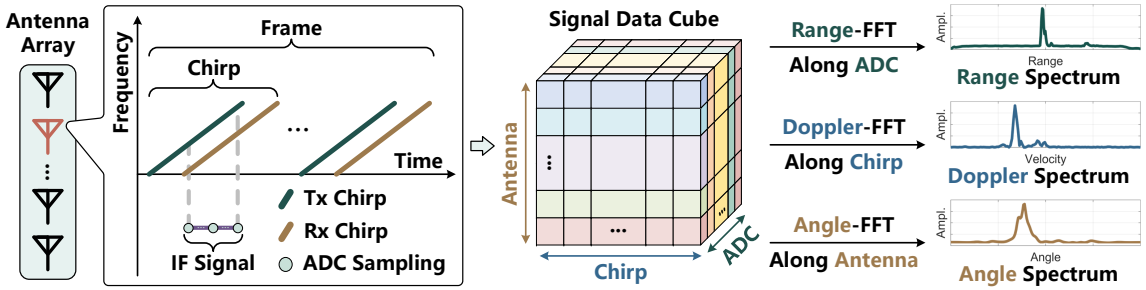


Fig. 2. mmWave signal model and spectra.

Challenge 3: How to derive identical secret offsets between the user and servers over an insecure channel? SpecRing requires the user and each authorized server to apply identical spectral offsets for encryption and decryption. However, directly transmitting the offsets or shared keys over the channel would immediately expose them to eavesdroppers. Moreover, using the same offsets for all servers would undermine fine-grained access control, since one server could infer dimensions that only another server is authorized to access. SpecRing addresses this challenge through pairwise key negotiation and deterministic parameter generation. The user negotiates an independent secret shared key with each authorized server without transmitting the key itself. For each radar frame, both sides combine this secret with a public salt to deterministically derive the same spectral offsets. The salt acts only as a public input and reveals no useful information without the shared key. Thus, eavesdroppers cannot infer the secret offsets from intercepted messages, while each server can derive only the offsets associated with its own authorization.

In summary, this paper makes the following contributions:

- We identify the hidden privacy risk of mmWave sensing despite its common privacy-friendly framing. We show that raw radar signals expose sensitive physical and physiological semantics, motivating protection mechanisms beyond conventional transport encryption.
- We propose a reversible semantic encryption mechanism that modulates raw mmWave signals with key-derived phase offsets to cyclically shift selected spectra. Unauthorized parties observe a semantically misleading virtual scene, whereas authorized services recover the true physical scene exactly.
- We design a dimension-selective access control mechanism for mmWave sensing. SpecRing allows each dimension to be independently encrypted and decrypted, enabling fine-grained authorization for multi-service ubiquitous sensing.
- We strengthen SpecRing against practical attacks. SpecRing eliminates environmental imprints and enables secure key negotiation and parameter generation over insecure channels, preventing adversaries with strong capabilities from inferring encryption parameters.

2 Preliminaries

2.1 Fundamentals of mmWave Signal Spectra

2.1.1 Signal Structure and Spectral Semantics. As shown in Fig. 2, an FMCW radar transmits chirps with initial frequency f_0 and sweep rate k. For a target at distance r, radial velocity v, and angle θ, the round-trip delay is $\tau = 2r/c$, where c is the speed of light. Mixing the transmitted and received signals yields the intermediate frequency (IF) signal

$$s(t) \approx A \exp\left(j2\pi\left(f_0\tau + k\tau t\right)\right), \tag{1}$$

where A is the signal amplitude. The IF signal encodes target information in the range, Doppler, and angle spectra. These spectra correspond to the ADC sample, chirp, and antenna dimensions of the radar signal cube, respectively. Suppose each chirp contains N_s samples acquired at rate f_s, each frame contains N_c chirps transmitted at repetition frequency f_c, and the radar uses a uniform linear array of N_a antenna elements with spacing d_a (typically $\lambda/2$), where $\lambda = c/f_0$. Let i_s, i_c, and i_a denote the sample, chirp, and antenna indices, respectively. The three spectra are obtained by applying FFTs along these dimensions:

$$\begin{aligned}
s(i_s) &= A\exp\left(j2\pi\left(f_0\tau + k\tau\frac{i_s}{f_s}\right)\right) \xrightarrow{\text{Range-FFT}} S_r(k_r) = \sum_{i_s=0}^{N_s-1} s(i_s)\exp\left(-j2\pi\frac{i_s k_r}{N_{f,r}}\right),\\
s(i_c) &= A\exp\left(j2\pi\left(f_0\tau + k\tau t + \frac{f_d}{f_c}i_c\right)\right) \xrightarrow{\text{Doppler-FFT}} S_d(k_d) = \sum_{i_c=0}^{N_c-1} s(i_c)\exp\left(-j2\pi\frac{i_c k_d}{N_{f,d}}\right),\\
s(i_a) &= A\exp\left(j2\pi\left(f_0\tau + k\tau t + \frac{d_a\sin\theta}{\lambda}i_a\right)\right) \xrightarrow{\text{Angle-FFT}} S_a(k_a) = \sum_{i_a=0}^{N_a-1} s(i_a)\exp\left(-j2\pi\frac{i_a k_a}{N_{f,a}}\right).
\end{aligned} \tag{2}$$

Here, $f_d = 2v/\lambda$ is the Doppler frequency, $N_{f,r}$, $N_{f,d}$, and $N_{f,a}$ are the FFT sizes, and k_r, k_d, and k_a are the corresponding spectrum-bin indices. The maximum unambiguous range, velocity, and angle are $r_{\max} = cf_s/(2k)$, $|v_{\max}| = \lambda f_c/4$, and $|\theta_{\max}| = \arcsin(\lambda/(2d_a))$, respectively. Combining the signal representations from the three dimensions, the IF signal can be written as

$$\begin{aligned}
s(i_s, i_c, i_a) &= A\exp\left(j2\pi f_0\tau + j2\pi k\tau\frac{i_s}{f_s} + j2\pi\frac{f_d}{f_c}i_c + j2\pi\frac{d_a\sin\theta}{\lambda}i_a\right)\\
&= A\exp(j2\pi f_0\tau)\cdot\exp\left(j2\pi k\tau\frac{i_s}{f_s}\right)\cdot\exp\left(j2\pi\frac{f_d}{f_c}i_c\right)\cdot\exp\left(j2\pi\frac{d_a\sin\theta}{\lambda}i_a\right).
\end{aligned} \tag{3}$$

Thus, the complete raw signal can be represented as a three-dimensional complex matrix $\mathbf{S} \in \mathbb{C}^{N_s\times N_c\times N_a}$.

The resulting range, Doppler, and angle spectra carry physical semantics at different scales. At the macroscopic level, range and angle peaks describe a target's spatial position, while Doppler peaks describe its radial velocity. At a finer-grained level, the spatial distribution and temporal phase variations provide signatures from which higher-level human information, such as posture, activity, and respiration, can be inferred. The factorization of Eq. 3 further shows that the three signal dimensions are mathematically separable.

2.1.2 Spectrum-Shifting Property. Spectrum-shifting is the process of changing the position of a signal within the frequency domain, which can be implemented by multiplying the time-domain signal by a complex carrier signal, as illustrated in Fig. 3. Suppose a time-frequency transform pair $s(i)\leftrightarrow S(k)$. Multiplying the time-domain signal $s(i)$ by the carrier signal $\exp(j2\pi\Delta k i/N)$ results in a cyclic offset of the frequency-domain representation $S(k)$ by Δk bins:

$$s(i)\cdot\exp\left(j2\pi\frac{\Delta k\cdot i}{N}\right) \xrightarrow{\text{FFT}} S\left((k-\Delta k)\bmod N\right), \tag{4}$$

where Δk is the frequency bin offset and N is the FFT size. Note that the modulo N operation reflects the inherent periodicity of the spectrum. When a frequency bin exceeds the maximum observable frequency, it wraps around to the beginning of the spectrum. Conceptually, if the two ends of the spectrum are connected, it forms a ring, and the spectrum-shifting process corresponds to a rotation along this ring, which inspires the name SpecRing.

Combining this property with the dimensional separability shown in Eq. 3, a key-controlled phase sequence applied along one signal dimension cyclically shifts its corresponding spectrum without shifting the other two. The radar cube retains its format and remains compatible with standard processing, while the observed bins

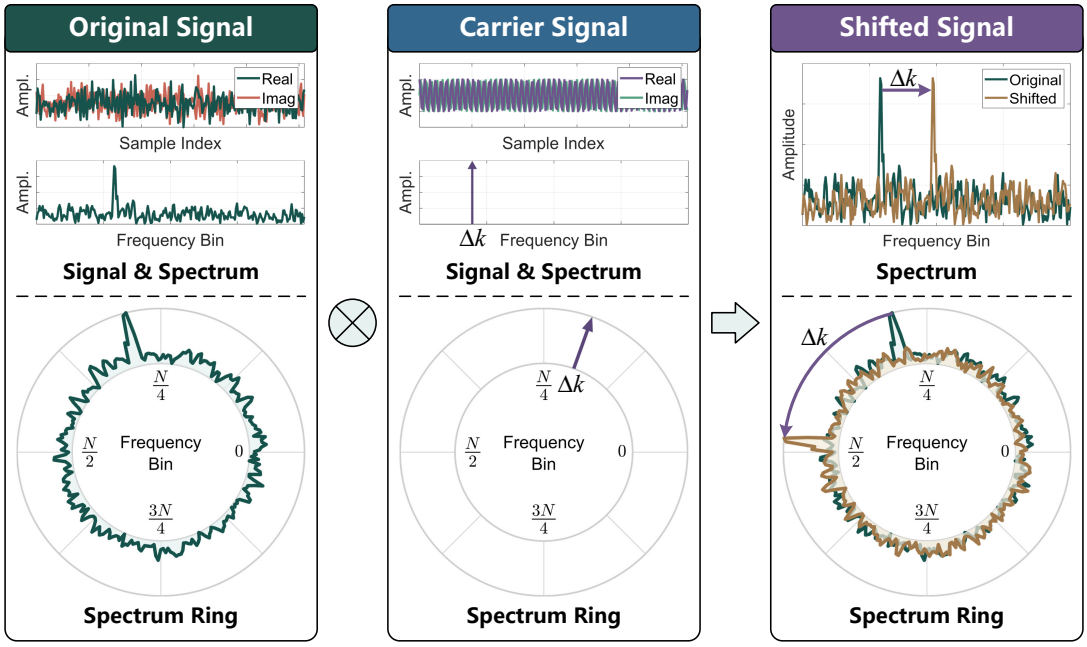


Fig. 3. Illustration of spectrum shifting.

no longer reveal the true physical state without the secret offsets. These properties provide the physical and mathematical basis for spectrum obfuscation, while keeping the offsets unknown requires key confidentiality.

2.2 Threat Model

We consider an outsourced sensing deployment in which the radar and SpecRing operate within a trusted local environment controlled by the user, while sensing tasks are performed by third-party servers. The user intentionally uploads radar-processable streams to these servers, and SpecRing protects each stream before it leaves the local environment. Within this deployment, we distinguish an honest-but-curious sensing server from an external eavesdropper.

2.2.1 Participants. The proposed SpecRing scheme involves the following three types of participants.

User: The owner of the mmWave radar and local SpecRing module. The user negotiates keys with sensing servers, collects raw signals, applies spectrum-shifting encryption, and distributes protected streams.

Sensing server: A third-party server that is authorized to perform an intended sensing task and is granted access to the required signal dimensions. It can decrypt the authorized spectra, but may analyze the received stream to infer information from dimensions outside its authorization.

External eavesdropper: Radar streams in outsourced sensing must traverse the user-server communication channel, creating a practical exposure point for eavesdropping. We therefore consider an unauthorized party that eavesdrops on this channel and obtains uploaded radar streams and key-negotiation messages.

Although the sensing server and external eavesdropper play different roles, their attack models are similar: both attempt to recover the true spectra in dimensions beyond their authorization. Thus, a sensing server is an adversary for its ungranted dimensions, whereas an external eavesdropper is an adversary for all dimensions. A

sensing server may also act as an external eavesdropper by monitoring communications between the user and other servers.

2.2.2 Adversary Capabilities. We consider a strong adversary with the following capabilities.

Algorithm knowledge: The adversary has complete knowledge of the proposed SpecRing scheme. All algorithmic details of spectrum-shifting encryption and key derivation are public.

Communication eavesdropping: The adversary can eavesdrop on all communications between the user and authorized servers, including the key negotiation messages.

Signal parameters: The adversary possesses all radar configuration and sampling parameters. In other words, the adversary has sufficient prior knowledge to perform correct and complete signal processing on any radar data available to it.

2.2.3 Security Goals. To defend against the adversary described above, SpecRing aims to achieve the following security guarantees.

Spectrum obfuscation: The spectrum-shifting encryption must hide the true spectral peak locations from unauthorized parties. When an adversary attempts decryption with an incorrect key, the resulting spectrum should exhibit peak locations that are indistinguishable from random guessing.

Dimensional independence: The encryption applied to one signal dimension must not compromise the information integrity of other dimensions. Each dimension should be independently encryptable and decryptable, enabling fine-grained, selective access control without cross-dimensional interference.

Key confidentiality: The key negotiation and parameter distribution mechanism must ensure that even if all transmitted messages are intercepted, the adversary cannot derive the shared keys used for spectrum encryption and decryption.

3 System Overview

SpecRing's encryption mechanism builds on the frequency-shifting property. At the system level, SpecRing integrates dimension-selective encryption, environmental imprint removal, and pairwise key negotiation into a privacy-preserving sensing framework, as illustrated in Fig. 4. For each authorized server, the user independently negotiates a shared key and deterministically derives per-frame spectral offsets. On the user side, a raw radar frame is organized into a three-dimensional signal cube, from which environmental imprints are removed to prevent adversaries from inferring secret offsets. The cleaned cube is subsequently encrypted through spectrum shifting using the derived offsets and distributed to authorized servers. Different servers may receive differently encrypted versions of the same radar stream. Each authorized server uses its pairwise key to recover only the permitted dimensions and performs standard sensing tasks without modifying existing pipelines. Unauthorized parties, in contrast, observe a format-valid radar stream whose spectral peaks are cyclically shifted to arbitrary locations, yielding a semantically misleading view of the physical environment while preserving the underlying signal structure.

4 Design of SpecRing

4.1 Spectrum-Shifting Encryption Algorithm

SpecRing's encryption is built on the mathematical foundation provided by the frequency-shifting property. By introducing phase offsets along each signal dimension, SpecRing cyclically rotates the corresponding spectra. Without knowledge of these offsets, an adversary cannot determine the true peak locations from the shifted spectrum, as all possible locations appear equally probable. The key holder can restore the spectrum to its original position by applying the inverse rotation. Let $s'(i) = s(i) \cdot \exp(j2\pi\Delta ki/N)$ denote the encrypted signal. The

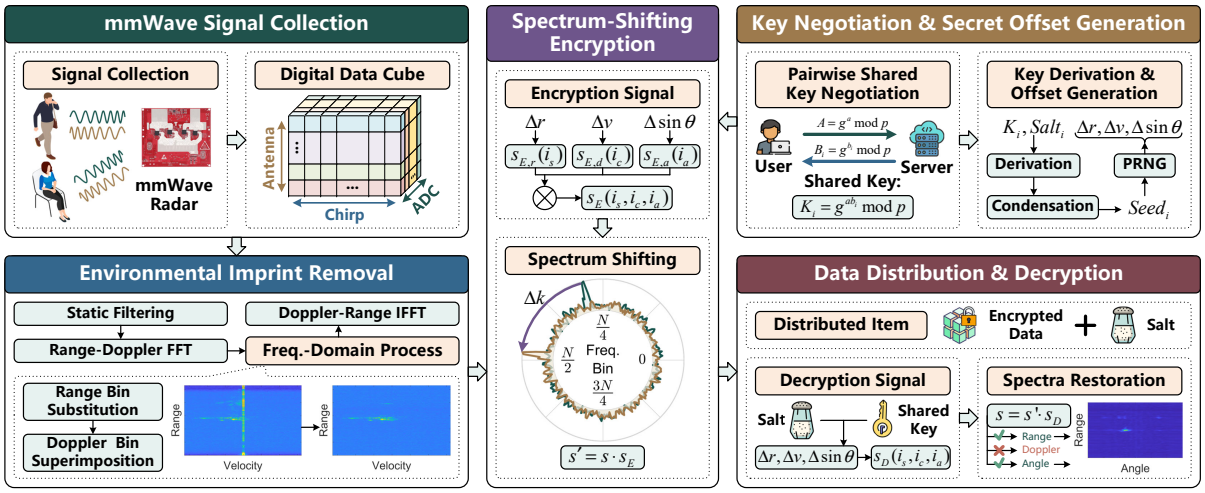


Fig. 4. System overview of SpecRing.

original signal can be recovered as

$$s(i) = s'(i) \cdot \exp\left(-j2\pi \frac{\Delta k \cdot i}{N}\right). \tag{5}$$

4.1.1 Encryption Process. The mmWave signal encodes target information across three orthogonal spectra: range, Doppler, and angle. SpecRing exploits the frequency-domain independence of these spectra to enable dimension-selective encryption. To facilitate a more intuitive description of this process, we rewrite the signal matrix from Eq. 3 as

$$s(i_s, i_c, i_a) = s_0 \cdot s_r(i_s) \cdot s_d(i_c) \cdot s_a(i_a), \tag{6}$$

where $s_0 = A\exp(j2\pi f_0 \tau)$ denotes the initial amplitude and phase, and

$$s_r(i_s) = \exp\left(\frac{j4\pi k i_s r}{c f_s}\right),\ s_d(i_c) = \exp\left(\frac{j4\pi i_c v}{\lambda f_c}\right),\ s_a(i_a) = \exp\left(\frac{j2\pi d_a i_a \sin\theta}{\lambda}\right), \tag{7}$$

represent the range-dependent, velocity-dependent, and angle-dependent components, respectively. The angular parameter is expressed as $\sin\theta$ rather than θ because the angular bins are uniformly spaced in the $\sin\theta$ domain.

Fig. 5 illustrates the overall spectrum-shifting encryption process. Here, the dimensional components are modeled as functions of the distance r, velocity v, and angle $\sin\theta$. Encrypting the true information is equivalent to introducing a distance offset Δr, velocity offset Δv, and angle offset $\Delta\sin\theta$. Thus, the encryption signal can be expressed as

$$s_E(i_s, i_c, i_a) = s_{E,r}(i_s) \cdot s_{E,d}(i_c) \cdot s_{E,a}(i_a), \tag{8}$$

where

$$s_{E,r}(i_s) = \exp\left(\frac{j4\pi k i_s \Delta r}{c f_s}\right),\ s_{E,d}(i_c) = \exp\left(\frac{j4\pi i_c \Delta v}{\lambda f_c}\right),\ s_{E,a}(i_a) = \exp\left(\frac{j2\pi d_a i_a \Delta\sin\theta}{\lambda}\right). \tag{9}$$

Due to the property that $\exp(ja)\cdot\exp(jb) = \exp(j(a+b))$, encryption is achieved by multiplying the original signal with the encryption signal, i.e.,

$$s'(i_s, i_c, i_a) = s(i_s, i_c, i_a) \cdot s_E(i_s, i_c, i_a) = s_0 \cdot \big(s_r(i_s) \cdot s_{E,r}(i_s)\big) \cdot \big(s_d(i_c) \cdot s_{E,d}(i_c)\big) \cdot \big(s_a(i_a) \cdot s_{E,a}(i_a)\big). \tag{10}$$

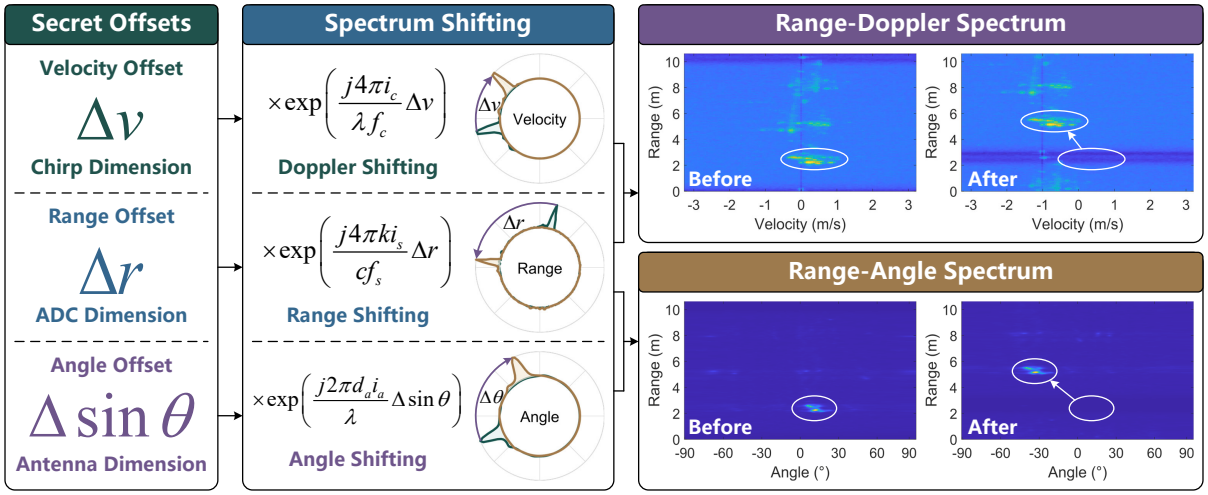


Fig. 5. The spectrum-shifting encryption process.

The encrypted distance, velocity, and angle are given by

$$r' = (r + \Delta r) \bmod r_{\max}, \tag{11}$$

$$v' = ((v + \Delta v + v_{\max}) \bmod 2v_{\max}) - v_{\max}, \tag{12}$$

$$\sin\theta' = ((\sin\theta + \Delta\sin\theta + \sin\theta_{\max}) \bmod 2\sin\theta_{\max}) - \sin\theta_{\max}. \tag{13}$$

Given secret offsets, the true spectral position of each target is cyclically shifted to a virtual location. We also demonstrate the effect of spectrum-shifting encryption on 2D spectra, including the range-Doppler spectrum (RDS) and the range-angle spectrum (RAS). Since the signal dimensions are mutually orthogonal, the offset for each dimension can be independently configured.

4.1.2 Decryption Process. The decryption process is the inverse of the encryption process. The decryption signal is constructed as the complex conjugate of the encryption signal, i.e.,

$$s_D(i_s, i_c, i_a) = s^*_{E,r}(i_s) \cdot s^*_{E,d}(i_c) \cdot s^*_{E,a}(i_a). \tag{14}$$

The original signal can be recovered by multiplying the encrypted signal with the decryption signal:

$$s(i_s, i_c, i_a) = s'(i_s, i_c, i_a) \cdot s_D(i_s, i_c, i_a). \tag{15}$$

Similarly, the target information is decrypted by

$$r = (r' - \Delta r) \bmod r_{\max}, \tag{16}$$

$$v = ((v' - \Delta v + v_{\max}) \bmod 2v_{\max}) - v_{\max}, \tag{17}$$

$$\sin\theta = ((\sin\theta' - \Delta\sin\theta + \sin\theta_{\max}) \bmod 2\sin\theta_{\max}) - \sin\theta_{\max}. \tag{18}$$

A key advantage of SpecRing is the ability to selectively encrypt individual dimensions while leaving others unencrypted. The dimensional signal components $s_r\,(i_s)$, $s_d\,(i_c)$, and $s_a\,(i_a)$ can be independently multiplied by their corresponding encryption signals $s_{E,r}\,(i_s)$, $s_{E,d}\,(i_c)$, and $s_{E,a}\,(i_a)$, respectively. Thus, servers can be granted access to different combinations of signal dimensions.

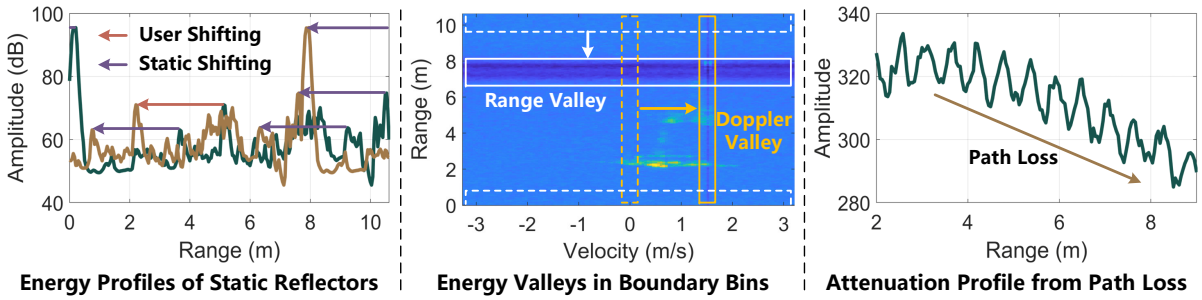


Fig. 6. Environmental imprints.

4.1.3 Discussion on Data Confidentiality and Utility. SpecRing protects privacy by altering physical semantics without destroying the signal structure required for sensing. For the two spatial dimensions, namely range and angle, spectrum shifting relocates a user from its true position to a virtual one, thereby concealing the user's real location from unauthorized inference. Meanwhile, SpecRing applies the same spectral shift to all bins within the encrypted dimension, so the complex signals contained in the user bins remain unaffected. A potential concern is that cyclic shifting may cause user bins to wrap around at the spectrum boundary. This does not fundamentally reduce utility, because the spectrum is inherently circular. A sensing system can easily address this by applying boundary extension before target detection.

For the Doppler dimension, spectrum shifting protects both coarse motion semantics and fine-grained physiological semantics. At the macroscopic level, it falsifies the target's apparent radial velocity. At the microscopic level, it misleads small-displacement measurements, thereby obscuring biometric information such as respiration and heartbeat. Meanwhile, as with the spatial dimensions, the signal structure within each Doppler bin is preserved, so target detection and signal extraction in the RDS space remain unaffected.

Overall, SpecRing changes what the data mean without breaking how the data are processed: unauthorized servers can still operate on a format-valid radar stream, but the recovered spatial and motion semantics correspond to a virtual scene rather than the true physical environment.

4.2 Environmental Imprint Removal

The spectrum-shifting encryption scheme described above possesses a critical vulnerability: while the target spectrum is shifted, the spectra of background components are shifted by the same offsets. These static background signatures serve as environmental imprints that can reveal the secret offsets to the adversary. As shown in Fig. 6, the environmental imprints comprise the following components.

Energy profiles of static reflectors: Environmental fixtures such as walls and furniture produce consistent reflection signals, forming distinct peaks in the spectrum. The shifted locations of these peaks directly reveal the secret offsets. Static targets also cause high energy in the zero-velocity bins of the RDS, resulting in a significant imprint in the Doppler spectrum.

Energy valleys in boundary bins: This type of imprint manifests in two aspects. On the one hand, range bins near zero distance have an extremely low noise level due to transmit leakage effects. After static filtering, these regions produce a distinct valley in the spectrum. Due to the cyclic nature of the spectrum, this valley also appears near the maximum unambiguous range. On the other hand, the Doppler dimension undergoes similar effects caused by static filtering. Since static targets possess zero Doppler frequency, the zero-velocity bins in the Doppler spectrum also form an energy valley after static filtering.

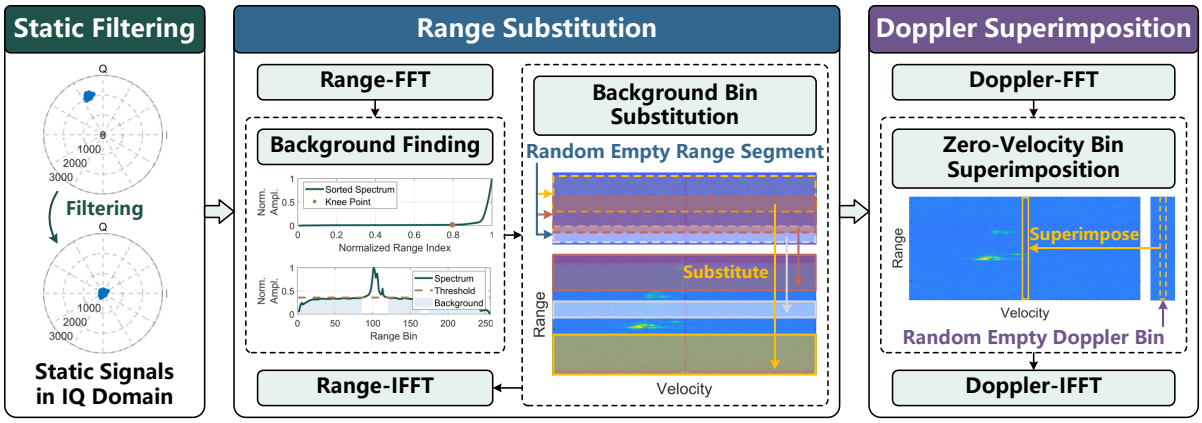


Fig. 7. The environmental imprint removal process.

Attenuation profile from path loss: Signal propagation suffers from distance-dependent path loss, creating a continuous decay pattern across the spectrum. This attenuation profile also serves as an exploitable imprint.

To address this issue, we propose an environmental imprint removal mechanism that eliminates the above signatures through three core operations: static target energy suppression, background range bin substitution, and zero-velocity Doppler bin compensation, thereby ensuring that an adversary cannot infer the secret offsets from the background spectra, as shown in Fig. 7.

4.2.1 Static Target Energy Suppression. We first address the spectral profiles of background reflectors, which can be eliminated through static filtering, a common technique in radar signal processing. Static targets in the environment, such as walls, furniture, and fixed infrastructure, produce consistent reflection signals across the chirps. In contrast, moving targets exhibit phase variations due to Doppler shifts. In the in-phase and quadrature (IQ) domain, static targets appear as compact clusters with minimal amplitude and phase variations, whereas dynamic targets manifest as dispersed point clouds with significant amplitude and phase fluctuations. This difference enables separation of static and dynamic components through temporal averaging. For each ADC sample index i_s and antenna i_a, we compute the mean signal across all N_c chirps in a frame:

$$\bar{s}(i_s, i_a) = \frac{1}{N_c} \sum_{i_c=0}^{N_c-1} s(i_s, i_c, i_a). \tag{19}$$

$\bar{s}(i_s, i_a)$ represents the cluster center of static signal points in the IQ domain. By subtracting it from the original signal, we effectively shift the center of static signal points to the origin in the IQ coordinate system:

$$s(i_s, i_c, i_a) \leftarrow s(i_s, i_c, i_a) - \bar{s}(i_s, i_a). \tag{20}$$

This operation suppresses the energy of static targets to the background clutter level, causing their corresponding peaks to vanish from the range spectrum. The same static filtering is also applied to the reference dataset $s_{ref}(i_s, i_c, i_a)$ to obtain a calibrated reference spectrum.

4.2.2 Background Range Bin Substitution. Even after static target suppression, boundary bins in the range dimension retain distinctive energy patterns that can reveal secret offsets. Specifically, range bins near zero distance and $r_{\max}$ form a distinct energy valley due to transmit leakage effects. Meanwhile, the path-loss-induced attenuation profile across the range spectrum remains as a continuous signature that can be exploited by

adversaries. To eliminate both the boundary valley and the path-loss attenuation profile, SpecRing substitutes the background range bins with segments drawn from unoccupied-room reference signals $s_{ref}(i_s, i_c, i_a)$. The reference signal captures the environmental response without human targets and can be updated manually or automatically whenever the environment is unoccupied.

We first apply range-FFT to obtain the range spectrum $S_r(k_r)$. To separate background bins from user bins, we sort the energy profile $|S_r(k_r)|$ and employ the Kneedle algorithm [32] on the sorted values. Let the sorted energy values be $|S'_r(1)| \le |S'_r(2)| \le \cdots \le |S'_r(N_r)|$, where N_r is the number of sorted bins. We normalize the indices and energy values to

$$x(i) = \frac{i}{N_r}, \quad i = 1, 2, \ldots, N_r, \tag{21}$$

$$y(i) = \frac{|S'_r(i)| - |S'_r(1)|}{|S'_r(N_r)| - |S'_r(1)|}, \tag{22}$$

and identify the knee point as the point with the maximum perpendicular distance to the diagonal line $y = x$:

$$i' = \arg\max_i \ \frac{|y(i) - x(i)|}{\sqrt{2}}. \tag{23}$$

The background bin set is then $\{k_r \mid |S_r(k_r)| \le |S'_r(i')|\}$, which is naturally partitioned into M contiguous segments. For each segment of length L_m, we randomly select a starting position in the reference range spectrum $S_{r,ref}$ and extract a segment of the same length. The extracted segment then replaces the corresponding background bins in $S_r(k_r)$. After substituting all background segments, we apply range-inverse FFT (IFFT) to transform the modified spectrum back to the time domain:

$$s(i_s) = \frac{1}{N_{f,r}} \sum_{k_r=0}^{N_{f,r}-1} S_r(k_r) \exp\left(j2\pi \frac{i_s k_r}{N_{f,r}}\right). \tag{24}$$

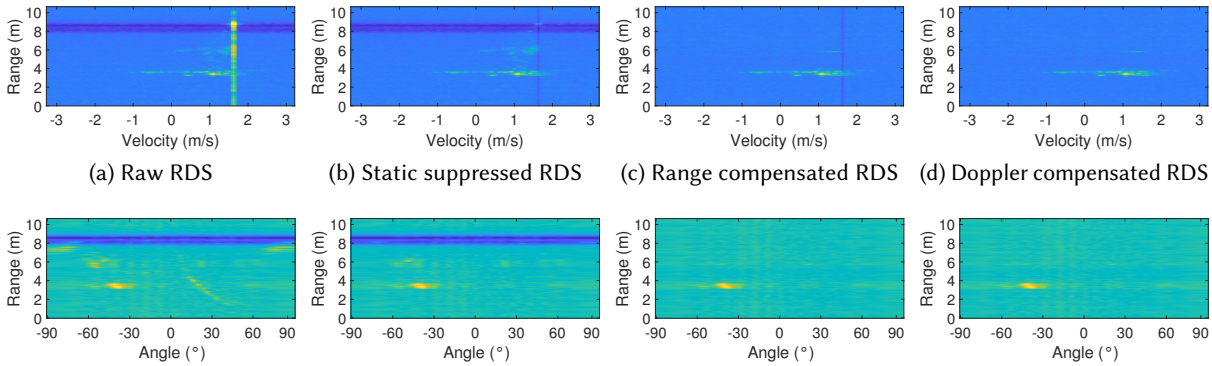
The reference spectrum shares similar environmental statistics (i.e., the noise floor and residual clutter) with the current frame, and therefore the substituted background appears natural. Furthermore, by randomizing the selection of reference segments, any stable spatial pattern such as the path-loss profile is broken.

Notably, dynamic objects confined to fixed regions, such as operating fans or pets in an enclosure, are not fully removed by static filtering and may also form stable anchors. SpecRing can include these persistent dynamic regions in the background bin set, either by identifying them through long-term energy monitoring or by allowing users to specify region masks, so that the signals in these regions can be replaced with the unoccupied-room reference signal. The implementation and evaluation in this paper focus on environments without such regions.

4.2.3 Zero-Velocity Doppler Bin Compensation. As with the boundary bins in the range spectrum, the zero-velocity bins in the Doppler spectrum form an identifiable energy valley after static filtering. Unlike range boundary bins, however, zero-velocity bins may contain valid signals from targets with near-zero radial velocity. Therefore, direct replacement of these bins could corrupt useful information. We observe that the energy level of the zero-velocity bin is significantly lower than that of ordinary background Doppler bins. Therefore, we compensate by superimposing a reference background bin onto the zero-velocity bin. This superposition raises the bin's energy to the background floor while keeping the resulting energy within the range of natural background variation.

Specifically, we first apply Doppler-FFT to obtain the Doppler spectrum $S_d(k_d)$. From the reference Doppler spectrum $S_{d,ref}$, we randomly select a non-zero-velocity bin, multiply it by a random phase offset $\phi_d \in_R [0, 2\pi)$, and superimpose the resulting signal onto the zero-velocity bin $S_d(k_0)$:

$$S_d(k_0) \leftarrow S_d(k_0) + S_{d,ref}(k_{bg}) \cdot \exp(j\phi_d), \tag{25}$$



(a) Raw RDS (b) Static suppressed RDS (c) Range compensated RDS (d) Doppler compensated RDS

(e) Raw RAS (f) Static suppressed RAS (g) Range compensated RAS (h) Doppler compensated RAS

Fig. 8. RDS (top row) and RAS (bottom row) at different stages of environmental imprint removal.

where k_{bg} is the index of the randomly selected non-zero-velocity bin. The random phase offset ϕ_d significantly increases the computational cost for an adversary who attempts to identify the superimposed bin through exhaustive search. Since the energy level of valid range-Doppler bins is significantly higher than that of background bins, the information in valid signals remains unaffected. Subsequently, we apply Doppler-IFFT to transform the compensated spectrum back to the time domain:

$$s(i_c) = \frac{1}{N_{f,d}} \sum_{k_d=0}^{N_{f,d}-1} S_d(k_d) \exp\left(j2\pi \frac{i_c k_d}{N_{f,d}}\right). \tag{26}$$

Fig. 8 illustrates the RDS and RAS at each stage of the environmental imprint removal process.

To support environmental imprint removal, SpecRing uses an unoccupied-room reference to replace background bins while maintaining a normal noise floor for downstream sensing. After static filtering shifts the IQ cluster centers of static reflectors to the origin, the reference contains only residual clutter and noise. Its noise floor may change slightly following major room-layout changes or long-term changes in radar hardware characteristics. Continuous radar collection allows SpecRing to refresh the reference whenever the room is detected to be unoccupied and free of dynamic objects. Each reference update requires only a short recording, even a single frame. This is because random segment selection during environmental imprint removal allows the same frame to be reused without revealing the secret offset. The reference can therefore be updated easily in practice.

4.3 Key Negotiation and Secret Offset Generation

SpecRing's encryption and decryption require a key negotiation process that addresses two challenges. First, the key negotiation must ensure key confidentiality. Even if all communications are intercepted, adversaries should be unable to derive the secret offsets. Second, fine-grained access control should be supported in multi-server environments, ensuring that each server can decrypt only the dimensions it is authorized to access and cannot infer the secret offsets of other servers. Fig. 9 shows the key negotiation and parameter generation.

4.3.1 User-Server Key Exchange. SpecRing utilizes the Diffie-Hellman key exchange (DHKE) protocol to negotiate pairwise keys between the user and each authorized server. DHKE enables two parties to establish a shared key over an insecure channel without directly transmitting the shared key itself. The protocol proceeds in three steps.

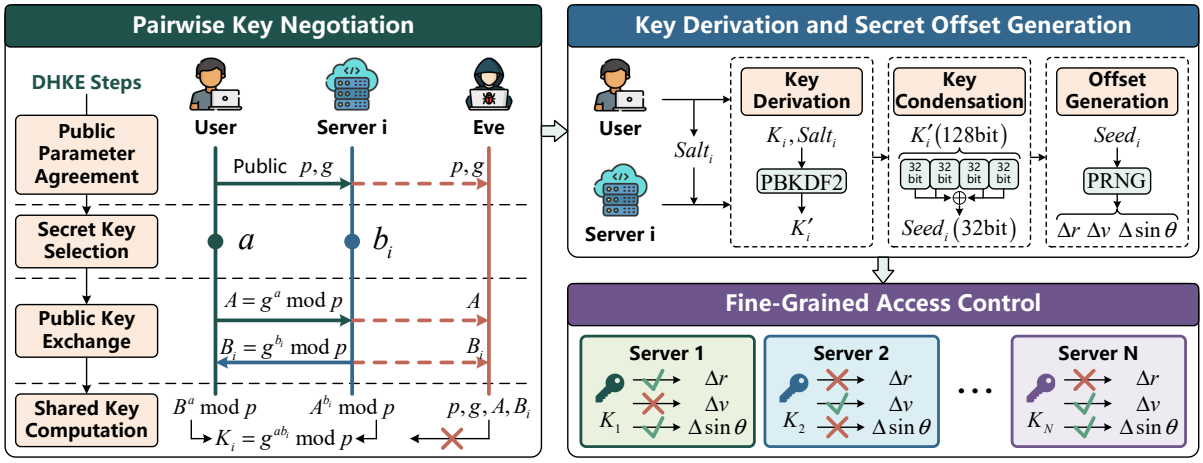


Fig. 9. Key negotiation and secret offset generation.

Public parameter agreement: First, the user and all servers publicly agree on a large prime number p and a small generator number g.

Private and public key generation: Then, the user generates a private key $a \in [2, p-2]$ and broadcasts the public key $A = g^a \bmod p$ to all servers. Concurrently, each server i generates its private key $b_i \in [2, p-2]$ and computes its public key $B_i = g^{b_i} \bmod p$, which is transmitted to the user.

Shared key computation: Finally, each server computes $A^{b_i} \bmod p$, while the user computes $B_i^a \bmod p$ for each server. Both parties arrive at the same shared key $K_i = g^{ab_i} \bmod p$.

The security of this key exchange stems from the discrete logarithm problem. Given $g^a \bmod p$ and $g^{b_i} \bmod p$, computing the shared key $K_i = g^{ab_i} \bmod p$ is equivalent to solving the discrete logarithm problem, which is a well-established computationally intractable problem. Consequently, even if an adversary intercepts all messages, including the public keys A and all B_i, they are still unable to compute the shared key K_i.

4.3.2 Key Derivation. The shared key K_i serves as the foundation for generating the secret offsets. However, using K_i directly to generate the secret offsets would be insecure because the resulting offsets would have limited entropy and be vulnerable to brute-force attacks. Therefore, SpecRing employs a key derivation process that derives a distinct key from the shared key for each frame, ensuring that the secret offsets for every frame are unique and secure. In each radar frame, the user generates a random 128-bit salt value $Salt_i$ for each server and then derives a 128-bit key K_i' using the password-based key derivation function 2 (PBKDF2) algorithm:

$$K_i' = \text{PBKDF2}(K_i, Salt_i), \tag{27}$$

which operates by iteratively applying a pseudorandom function to the shared key and salt. The pseudorandom function is typically instantiated as the hash-based message authentication code algorithm, which internally employs a cryptographic hash function such as SHA-256. In each iteration round, the function takes the shared key and salt as input, produces an output value, and feeds this output into the next iteration. The final output is extracted as the derived key. The salt value for each frame is transmitted to the server along with the data, enabling the server to compute the same derived key.

4.3.3 Secret Offset Generation. Finally, we describe the generation of secret offsets from the derived key. These offsets represent random shifts that should be uniformly distributed over one period of the spectral ring. Given

that the 128-bit derived key exceeds the required entropy for this application, we condense it by partitioning it into four 32-bit segments and applying bitwise exclusive OR (XOR) to produce a 32-bit random seed:

$$Seed_i = \bigoplus_{j=0}^{3} K_i'[j \cdot 32 : (j+1) \cdot 32 - 1], \tag{28}$$

where $\bigoplus$ denotes the bitwise XOR operation. This seed is then used to initialize a pseudorandom number generator (PRNG) for producing the secret offsets:

$$\Delta r_i = \text{PRNG.next}() \in_R [0, r_{\max}), \tag{29}$$

$$\Delta v_i = \text{PRNG.next}() \in_R [-v_{\max}, v_{\max}), \tag{30}$$

$$\Delta \sin\theta_i = \text{PRNG.next}() \in_R [-\sin\theta_{\max}, \sin\theta_{\max}). \tag{31}$$

The server derives the same random seed from the shared key and initializes the same PRNG to generate the same secret offsets.

Since each dimension is encrypted independently and each server negotiates a distinct shared key, SpecRing naturally supports dimension-selective access control. The user can selectively encrypt different dimensions for different servers using their respective secret offsets, and transmit the individually encrypted signals to each server separately. Each server derives the secret offsets for its granted dimensions from its pairwise shared key and restores the corresponding spectra. Dimensions not granted to a server remain encrypted in the signal it receives, since their offsets are derived from a different shared key.

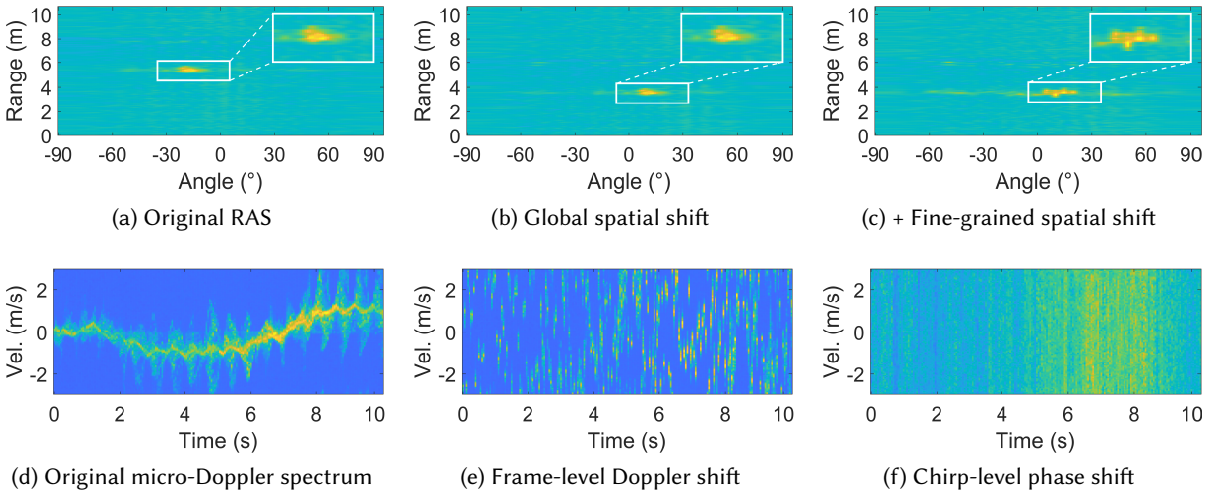
4.3.4 Discussion on Fine-Grained Access Control. SpecRing currently supports task-specific combinations of signal dimensions. However, global spectrum shifting mainly conceals macroscopic physical semantics, while local spectral patterns may still reveal fine-grained human information. For example, shifting a person's apparent position does not conceal their posture. We next discuss how spectrum shifts can be applied at finer granularities to support more selective authorization.

For spatial dimensions, FFT processing first transforms the raw signal cube into a range-angle cube. The angle spectrum at each range bin and the range spectrum at each angle bin can then be cyclically shifted using small, independent offsets. These offsets can also be derived from the shared key, allowing the user and an authorized server to generate the same offsets. These fine-grained shifts can disturb shape-related spatial patterns while retaining the target's coarse position, as shown in Fig. 10c. Such shifts can also be confined to selected range-angle regions. With separate keys for global and fine-grained shifts, a sensing server can be authorized to recover only the spatial information required by its task, such as a coarse position, information from selected regions, or fine-grained shape information.

For temporal protection, SpecRing derives a distinct per-frame key and corresponding spectral offsets from the shared key. These frame-wise offsets break cross-frame continuity in sensing features that encode temporal information, as shown by the encrypted micro-Doppler spectrum in Fig. 10e. At a finer scale, each chirp can receive a distinct phase offset, disrupting inter-chirp phase coherence and Doppler structure, as shown in Fig. 10f. Since the offset is shared across all ADC samples and antenna channels within a chirp, the range and angle information remains unaffected. The later case study on respiration monitoring will demonstrate how these two temporal granularities disrupt respiration waveform extraction.

4.4 Discussion on Protection Properties

To clarify how SpecRing differs from other protection approaches, we compare it with conventional encryption, random phase perturbation, and MuFilter [23]. Conventional encryption transforms the serialized radar stream into ciphertext, while phase perturbation applies a key-derived phase offset to each complex sample. MuFilter



(a) Original RAS (b) Global spatial shift (c) + Fine-grained spatial shift

(d) Original micro-Doppler spectrum (e) Frame-level Doppler shift (f) Chirp-level phase shift

Fig. 10. Fine-grained spatial and Doppler spectrum perturbations.

selectively perturbs behavioral signatures while retaining the signal cube format. Their protection properties are summarized in Table 1.

Table 1. Functional comparison of protection methods.

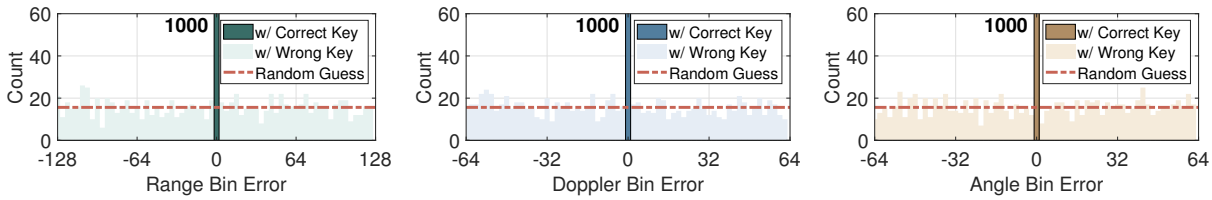
Property	Conventional encryption	Phase perturbation	MuFilter [23]	**SpecRing**
Protection output	Ciphertext	Distorted spectra	Sanitized spectra	**Shifted spectra**
Signal cube format retained	No	Yes	Yes	**Yes**
Dimension-selective protection	No	No	Yes	**Yes**
Resistance to eavesdropping	Yes	Yes	No	**Yes**

The comparison highlights key differences in their protection properties. Conventional encryption and phase perturbation protect the radar stream as a whole and therefore cannot provide selective access to individual signal dimensions. MuFilter supports dimension-selective sanitization while retaining the signal cube format, but its unencrypted output allows channel eavesdroppers to infer behavioral information that is not explicitly sanitized. In contrast, SpecRing provides both dimension-selective protection and resistance to channel eavesdropping. It retains a radar-compatible signal cube, allows each server to recover only its authorized dimensions, and prevents external eavesdroppers from obtaining the true spectra.

5 Basic Evaluation

5.1 Experimental Setup

We first evaluate SpecRing's core encryption algorithm in terms of its confidentiality, decryption reversibility, dimensional isolation, and key sensitivity. We also conduct ablation experiments on the environmental imprint removal mechanism and measure SpecRing's computational overhead. The experimental data are collected from the same indoor human-tracking environment as Case Study 1. We use a TI MMWCAS-RF-EVM radar platform with an 11-wavelength virtual array configuration. The effective bandwidth is 3.6 GHz, the chirp period is 300 μs,



(a) Range (b) Doppler (c) Angle

Fig. 11. Error distribution in encryption effectiveness evaluation.

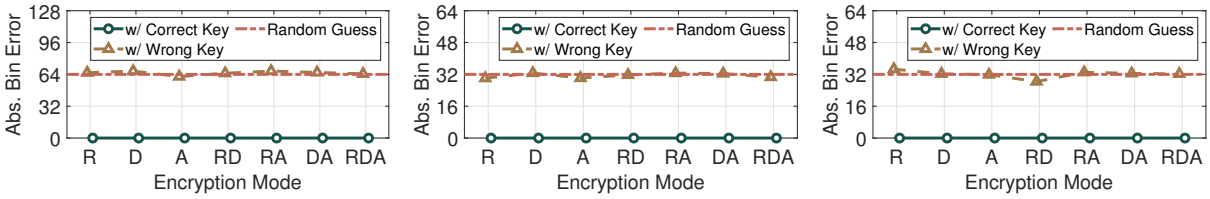
and the frame period is 50 ms. Each frame consists of 128 chirps with 256 ADC samples per chirp. Most evaluations contain 1000 trials, and most trials use a single frame of single-person tracking, randomly selected from 20000 frames collected across 100 tracking records. This dataset provides clear spectral peaks in range, angle, and velocity dimensions, and is well suited for the evaluations of spectrum-shifting encryption. The FFT sizes are 256 for range, 128 for Doppler, and 128 for angle. All computations are implemented in MATLAB R2024b and executed on a PC with an Intel i9-14900K CPU and 32 GB RAM. All experiments were conducted in accordance with our institute's standard institutional review board (IRB) protocol.

5.2 Evaluation on Encryption Effectiveness

We evaluate whether spectrum-shifting encryption successfully conceals the true spectral peak locations from unauthorized parties while allowing authorized parties to recover them. In each trial, we compute a new shared key, derive secret offsets for the three dimensions, and encrypt each spectrum through cyclic shifting. We then measure the relative bin error between the decrypted and original peak indices under both the correct shared key and a randomly generated key (labeled as the wrong-key condition). Since the spectrum is circular, the error is computed as the relative cyclic distance. For an N-bin spectrum, the error lies in $[-N/2, N/2)$. The results are shown in Fig. 11. With the correct key, the errors for all dimensions are zero, confirming that authorized decryption completely recovers the original spectra. Under the wrong-key and no-key conditions, the errors are uniformly distributed across the cyclic spectrum, consistent with random guessing. The adversary gains no information about the true peak locations, demonstrating that spectrum-shifting encryption effectively conceals spectral semantics from unauthorized parties.

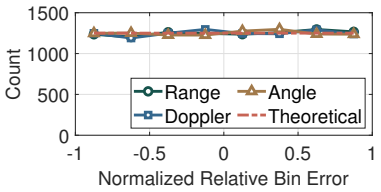
5.3 Evaluation on Dimension-Selective Access Control

SpecRing allows the user to selectively encrypt different dimensions, enabling fine-grained access control. To evaluate this, we enumerate all seven encryption combinations in total: R (range only), D (Doppler only), A (angle only), RD, RA, DA, and RDA. For each combination, we encrypt only the specified dimensions, while leaving the remaining dimensions unencrypted. We measure the absolute bin error under decryption with the correct key and decryption with a wrong key. The no-key condition simulates an attacker who is unaware that the signal is encrypted and attempts to obtain peak locations directly. For an N-bin spectrum, the absolute error is capped at $N/2$, while random guessing yields an expected error of $N/4$. Fig. 12 shows the per-dimension absolute bin error for each encryption combination. Correct-key decryption completely recovers peak positions for authorized dimensions, whereas wrong-key decryption yields errors consistent with random guessing, demonstrating SpecRing's confidentiality and reversibility under multi-dimensional encryption. The dimensions are well isolated, and an authorized server can access semantic information only for its granted dimensions without inferring offsets for others.

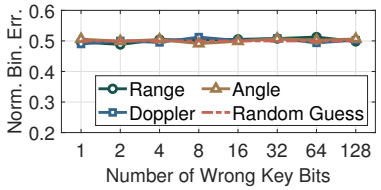


(a) Range (b) Doppler (c) Angle

Fig. 12. Effect of dimension-selective encryption modes.



(a) One-bit key error



(b) Multi-bit key errors

Fig. 13. Results on key sensitivity evaluations.

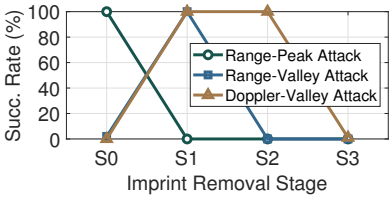


Fig. 14. Ablation study of environmental imprint removal.

5.4 Evaluation on Key Sensitivity

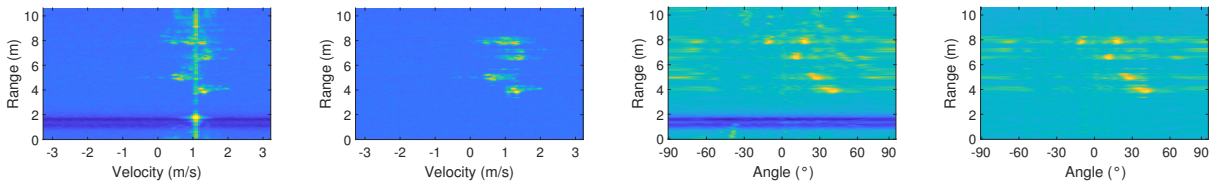
A secure encryption scheme must exhibit a strong avalanche effect: a tiny change in the key should produce completely unrelated ciphertext (or, in our case, completely unrelated offsets). We evaluate key sensitivity from two perspectives. First, we evaluate the normalized error distribution when the key deviates from the correct one by a single bit. As shown in Fig. 13a, a one-bit error produces normalized errors indistinguishable from the theoretical distribution, meaning the decrypted offsets are statistically independent of the correct ones. This confirms that SpecRing exhibits a strong avalanche effect, preventing the adversary from converging toward the correct key through the decrypted result trends.

Then, we introduce $2^0, 2^1, \ldots, 2^7$ error bits into the shared key, and measure the normalized bin errors across all dimensions. Fig. 13b shows that the bin error remains at the random-guessing level across all scales of key error. Together, these two evaluations confirm that an adversary gains no partial information from wrong-key decryption results and thus cannot converge toward the correct key, ensuring SpecRing's strong resistance to brute-force attacks.

5.5 Evaluation on Environmental Imprint Removal

Environmental imprints can serve as anchors for an adversary to estimate the secret offsets even without the key. We evaluate the effectiveness of our three-stage imprint removal mechanism through an ablation study with four configurations. S0 directly applies spectrum shifting to the raw signal without imprint removal, serving as the simple frequency-shift baseline. S1 applies static filtering on top of S0, S2 adds range compensation to S1, and S3 further adds Doppler compensation to complete the three-stage pipeline.

We implement three attack strategies that exploit environmental imprints to estimate offsets. The range-peak attack locates the highest range peak (corresponding to the high-energy peak near zero distance caused by transmit leakage) to infer the range offset. The range-valley attack finds the minimum-energy valley left by static



(a) RDS w/o Removal (b) RDS w/ Removal (c) RAS w/o Removal (d) RAS w/ Removal

Fig. 15. Encrypted RDS and RAS of a five-person signal before and after environmental imprint removal.

filtering to estimate the range offset. The Doppler-valley attack finds the minimum-energy Doppler bin in each range bin, thereby revealing the zero-velocity line and hence the Doppler offset. We define attack success as the case where the absolute cyclic bin error between the estimated offset and the true offset is ≤ 2 bins.

This evaluation includes multi-person tracking data, which increases the challenge of imprint removal. Fig. 15 compares the RDS and RAS of a five-person signal before and after imprint removal. Before removal, the RDS exhibits a distinct near-zero range peak, boundary range valley, and zero-velocity line, while the RAS shows a distinct range valley. All of these provide reliable anchors for attacks. After removal, these anchors are effectively suppressed, substantially raising the attack difficulty.

Fig. 14 reports the attack success rates in detail. Under S0, the range-peak attack achieves 100% success because direct spectrum shifting preserves the transmit-leakage peak as an environmental anchor. Static filtering successfully removes anchors for range-peak attacks, driving their success rate to the random-guessing level (below 1%), but range-valley and Doppler-valley attacks become effective because filtering introduces valleys. After adding range compensation, the range-valley attack success rate drops below 1%. Finally, Doppler compensation removes the zero-velocity line, reducing the Doppler-valley attack success rate to below 1%. With full imprint removal, all three attacks fall to the random-guessing level, demonstrating the mechanism's effectiveness and necessity.

5.6 Computational Cost

SpecRing requires no model training or inference. Its per-frame processing mainly consists of element-wise complex multiplications and FFT/IFFT operations. We record the per-frame computational overhead of each SpecRing component, and the results are shown in Fig. 16. Key derivation takes 7.1 ms. The range and Doppler encryption are fast, requiring only 1.0 ms. Angle encryption incurs additional overhead due to multiple-input multiple-output (MIMO) processing, but still completes in only 1.9 ms. Encrypting all three dimensions jointly takes 3.1 ms in total. The environmental imprint removal process takes 22.5 ms in total. Decryption takes 3.2 ms, similar to full encryption. The processing time is 25.6 ms on the user side and 10.3 ms on the server side, demonstrating that SpecRing introduces low computational overhead and meets real-time requirements. Key derivation can be completed during signal collection and is therefore excluded from the user-side total time. In addition, one DHKE execution takes 3.3 ms and is not included in the per-frame processing time.

6 Case Study

We evaluate SpecRing across four radar platforms and six sensing tasks, as summarized in Fig. 17. Since the basic evaluation has established decryption reversibility and dimensional isolation, these case studies focus on whether SpecRing blocks unauthorized sensing of privacy-sensitive semantics. We further evaluate adaptive inference attacks that combine learning models with prior knowledge.

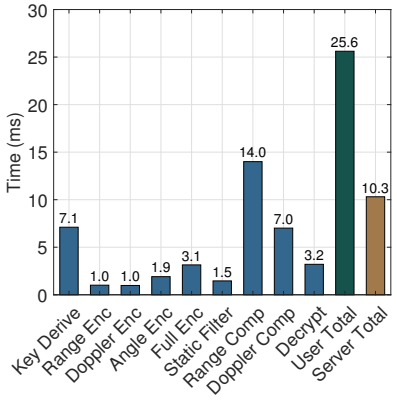


Fig. 16. Computational time per frame.

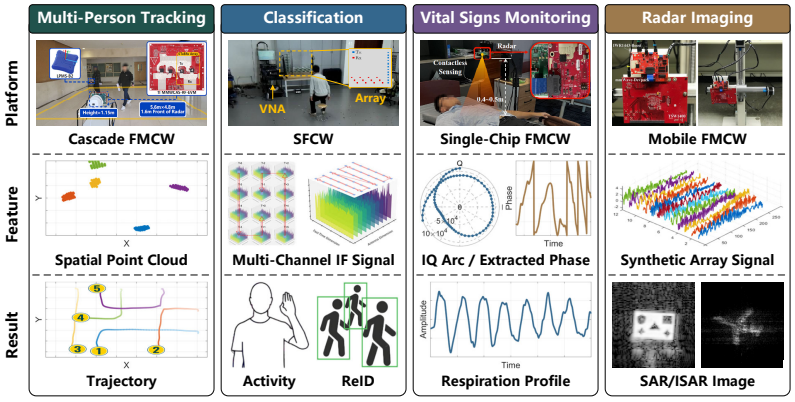
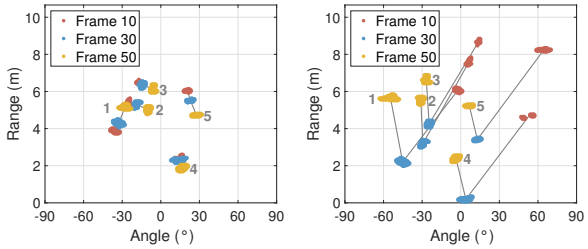
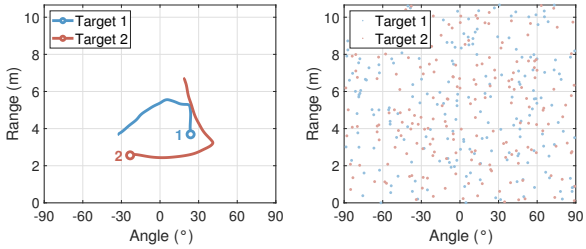


Fig. 17. Case study overview.



(a) Original (b) Encrypted

Fig. 18. Example of inter-frame cluster association.



(a) Original (b) Encrypted

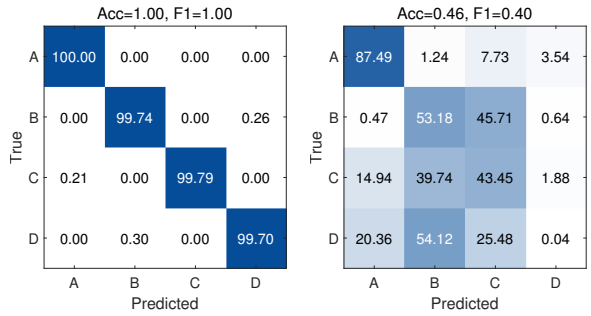
Fig. 19. Target tracking results before and after encryption.

6.1 Multi-Person Tracking

To assess the protection of absolute positions, we use the data collection setup and tracking pipeline from PmTrack [20]. The pipeline extracts a point cloud from the RAS through constant false alarm rate (CFAR) detection. Before CFAR, we pad the RAS to prevent cyclically shifted clusters from being split at spectrum boundaries. We encrypt the range and angle dimensions and process the encrypted signal through the same pipeline. For five users observed over two seconds, Fig. 18 shows that the original clusters remain close across frames, whereas encryption preserves their relative configuration but displaces them by large frame-specific offsets, breaking cross-frame association. Fig. 19 shows that the encrypted coordinates of two targets are scattered and no longer correspond to their true trajectories, thereby concealing the users' absolute positions.

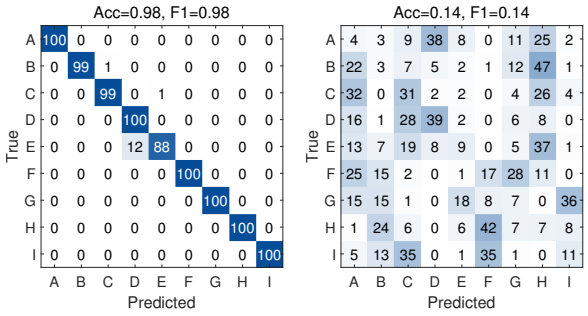
6.2 Classification

The classification case considers activity recognition and user re-identification (ReID), which target behavioral and identity information, respectively. We adopt the RFMamba [51] framework for data collection and algorithm deployment. The platform is a stepped-frequency continuous-wave (SFCW) radar driven by a vector network analyzer, with a 12×12 antenna-channel configuration and 201 stepped-frequency points. RFMamba directly takes the IF signal stream as input and learns discriminative spatiotemporal signatures end to end without relying on handcrafted sensing features. The classifier is trained on the original unencrypted signal stream and kept fixed



(a) Original (b) Encrypted

Fig. 20. Activity classification performance.



(a) Original (b) Encrypted

Fig. 21. Re-ID performance.

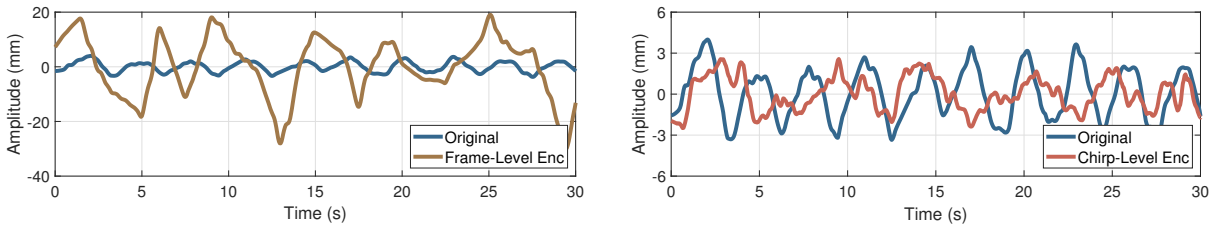
during the privacy evaluation. We encrypt both the spatial and Doppler dimensions and feed the encrypted stream into the same unauthorized pipeline. This setup models a zero-shot transfer attack in which an unauthorized server directly applies a well-trained model learned from unencrypted sensing data to encrypted inputs without knowing the secret offsets. The dataset contains measurements from 9 users performing four daily activities (walk, wave, sit, and static pose) over about 8 hours; we report row-normalized confusion matrices, accuracy, and F1-score.

Activity recognition: For the four activities denoted A–D, Fig. 20 shows nearly perfect recognition on the original measurements, whereas encryption reduces the accuracy and F1-score to 0.46 and 0.40, respectively. The accuracy remains above the random-guessing level of 0.25 because recognition is highly uneven across activities. Walking remains comparatively recognizable because its whole-body motion is more distinct, whereas the disturbed spatial and Doppler spectra make waving and sitting difficult to distinguish and cause static-pose samples to be misclassified as dynamic activities.

User re-identification: For the 9 users denoted A–I, Fig. 21 shows that the accuracy and F1-score both drop from 0.98 to 0.14, close to the random-guessing level ($1/9 \approx 0.11$). This result indicates that individual-specific signatures are effectively obfuscated.

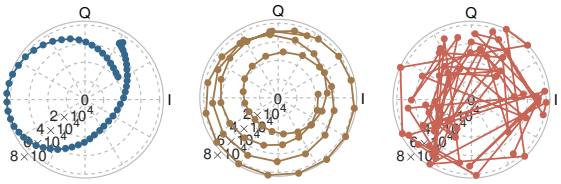
6.3 Vital Signs Monitoring

For respiration monitoring, we use the scenario and radar configuration in [6], with a single-chip FMCW radar operating at a 200 Hz chirp rate. Respiration waveforms are obtained using the RF-GymCare [38] algorithm, which follows a standard phase-extraction pipeline. We apply Doppler encryption at two granularities. Frame-level encryption groups 200 consecutive chirps into a one-second frame and applies a single random Doppler offset to the group. Chirp-level encryption instead applies an independent random phase offset to each chirp. Fig. 22 shows that both granularities severely distort the respiration waveforms, although the underlying mechanisms differ. The side of a Doppler bin relative to the zero-Doppler bin determines the direction of IQ rotation, while its distance from zero determines the angular speed. Fig. 23 shows a 60-sample IQ arc segment. The original IQ trace follows a counter-clockwise arc, and the extracted phase grows quasi-linearly over roughly one full cycle, as shown in Fig. 24. However, after frame-level encryption, the Doppler component moves to the opposite side of the zero-Doppler bin and farther from it. The IQ arc therefore rotates clockwise at a higher rate, causing the extracted phase to drop rapidly. This distortion is unpredictable and varies from frame to frame, severely corrupting the recovered respiration waveform. Chirp-level encryption disrupts the inter-chirp phase coherence more drastically, turning both the signal phase and the respiration waveform into chaotic random fluctuations.



(a) Frame-level (b) Chirp-level

Fig. 22. Respiration waveforms under frame-level and chirp-level encryption.



(a) Original (b) Frame-level Enc (c) Chirp-level Enc

Fig. 23. An IQ segment of respiration.

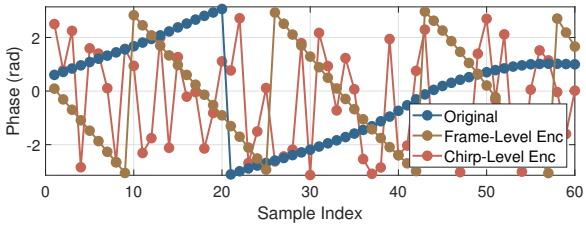


Fig. 24. Extracted phase segment.

In summary, Doppler encryption does not merely relocate spectral components. It disrupts the temporal phase continuity required for respiration recovery, preventing unauthorized servers from extracting a usable waveform whether chirps are uploaded individually or batched into frames.

6.4 Radar Imaging

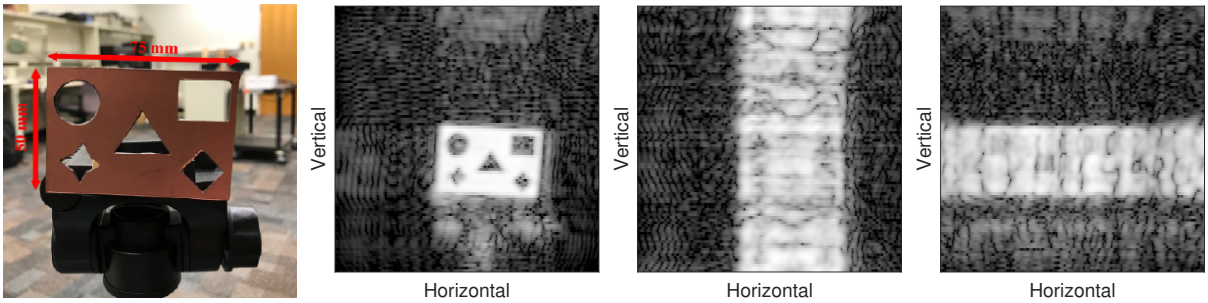
The imaging case examines whether SpecRing conceals object geometry in two radar-imaging tasks. Near-field synthetic aperture radar (SAR) uses array and synthetic-aperture dimensions, whereas far-field inverse synthetic aperture radar (ISAR) forms object profiles from range-Doppler measurements.

SAR imaging: We use the dataset and algorithm from [47], in which an IWR1443 Boost radar provides the vertical array and is scanned along the horizontal axis to synthesize a horizontal aperture. The original image in Fig. 25b preserves the target's major structural features. Encrypting the vertical array produces a vertical stripe, while encrypting the synthetic aperture produces a horizontal smear. In both cases, the reconstructed image no longer reveals the true geometry.

ISAR imaging: We use simulated radar signals of the Yak-42 aircraft and the reconstruction algorithm from CVAE [37]. The original image in Fig. 26b shows a recognizable aircraft profile with a fuselage and wing-like structures. Since ISAR azimuth is coupled to range through Doppler-dependent motion, we test Doppler-only and joint range-Doppler encryption. Doppler-only encryption collapses the aircraft into range-dependent horizontal patterns, while joint encryption produces a noise-like image with no identifiable structure.

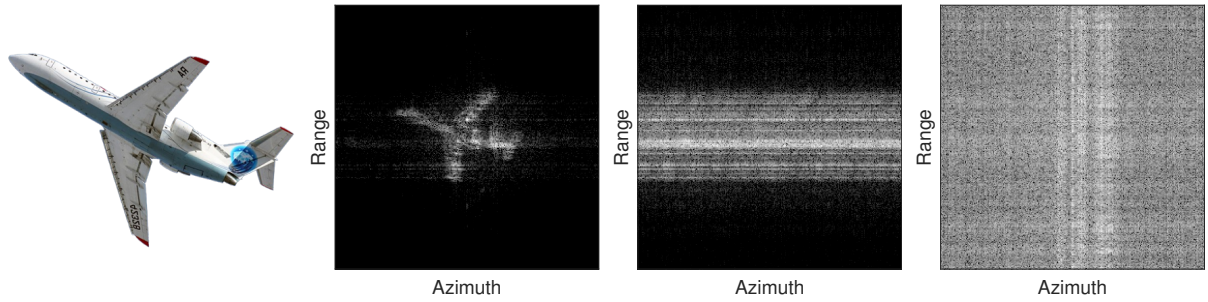
6.5 Adaptive Inference Attacks

Adaptive inference attacks are evaluated from two attacker perspectives. An external eavesdropper attempts to recover the target's physical states using layout and motion priors, while an authorized server attempts to infer encrypted dimensions from partial plaintext information. All plaintext-trained models achieve strong performance



(a) Photo reference (b) Original (c) Array encryption (d) Aperture encryption

Fig. 25. SAR imaging performance under different encrypted spatial dimensions.



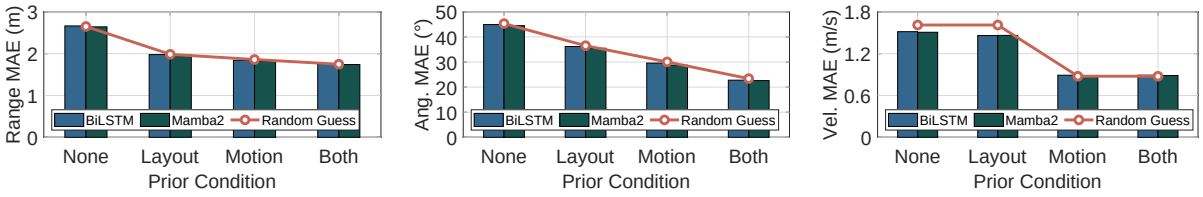
(a) Photo reference (b) Original (c) Doppler encryption (d) Range-Doppler encryption

Fig. 26. ISAR imaging performance under different encrypted dimensions.

on plaintext inputs. Accordingly, plaintext inputs are evaluated with these models, whereas encrypted inputs are evaluated with the corresponding models fine-tuned on encrypted samples.

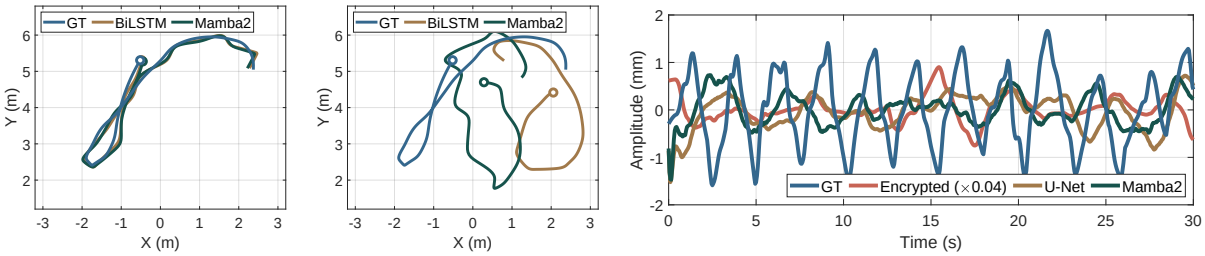
Inference with prior knowledge: We pair Swin V2 [25] with BiLSTM [13] or Mamba2 [9] to estimate range, angle, and radial velocity from multiframe RAS and RDS sequences. We represent room layout knowledge as a feasible walking area encoded in an additional RAS mask channel, while motion constraints are incorporated through temporal losses and bound penalties. After encryption, Fig. 27 shows that their errors remain close to the random guess baseline under all prior knowledge conditions. Thus, prior knowledge narrows the feasible state space but does not reveal the independently refreshed offsets.

Inference with partial plaintext: Some information leakage across dimensions is unavoidable. Successive range measurements reveal radial velocity, while integrating radial velocity yields displacement. Beyond these explicit relationships, we evaluate whether partial plaintext exposes any additional protected information. We consider two inference directions, from plaintext Doppler spectra to protected position and from plaintext position to fine-grained Doppler features. For position inference, we encrypt range and angle while leaving Doppler in plaintext, and evaluate whether the models can recover position. As shown in Fig. 28b, the position errors appear moderate mainly because these priors confine the predicted trajectories to the limited walking area. The models capture some relative range variation from plaintext Doppler but do not recover absolute range or angle, and the result therefore does not indicate successful spatial inference. For inference of fine-grained Doppler features,



(a) Range (b) Angle (c) Radial velocity

Fig. 27. Physical state recovery errors under different prior knowledge conditions.



(a) Plaintext position (b) Encrypted position (c) Encrypted respiration

Fig. 28. Server-side position and respiration inference.

a 1D U-Net [30] and a variant with a bidirectional Mamba2 [9] bottleneck reconstruct respiration waveforms from normalized IQ sequences at the correct chest range bin. As shown in Fig. 28c, the reconstructions from encrypted inputs no longer follow the respiration cycles, indicating that plaintext range does not enable recovery of fine-grained Doppler features.

7 Related Work

7.1 mmWave-based Human Sensing

mmWave sensing has become an important modality for contactless human sensing, with recent work enabling increasingly fine-grained spatial measurements [17, 45]. Existing studies have demonstrated mmWave-based multi-person tracking [20], activity, gesture, and hand-motion recognition [5, 11, 22, 28, 55], user identification [46], and human shape reconstruction [4, 7]. These capabilities allow mmWave systems to infer rich macroscopic behavioral information, such as user location, posture, motion, and body shape. Beyond macroscopic behavior sensing, mmWave signals can also capture fine-grained micro-motions [44]. This sensitivity supports speech enhancement [40] and health-related sensing, including respiration monitoring [2, 36, 50], heartbeat and cardiac activity sensing [48, 49], blood pressure estimation [14], and long-term health monitoring [3]. Together, these studies show that radio sensing can reveal not only external behaviors but also physiological states. However, most sensing systems assume trusted downstream processing, leaving raw radar streams vulnerable to misuse beyond the intended sensing task.

7.2 mmWave Sensing Privacy and Security

Existing studies have investigated mmWave sensing privacy and security from the physical layer to the data layer. At the physical layer, crafted signals [15, 35] or engineered reflectors [8, 29, 34] can spoof radar observations by creating, concealing, or displacing perceived targets. These methods may degrade the underlying sensing signal and affect authorized sensing services using the same data. Malicious privacy inference exploits mmWave radar's sensitivity to subtle motion, enabling the recovery of speech content [19, 39, 43, 54] or phone-call content [1] from sound-induced vibrations, as well as data extraction from vibrating storage devices [42]. Such eavesdropping can even occur through walls and reveal diverse private information [18, 27, 52].

At the data level, recent studies have explored removing sensitive information from sensing data while preserving its utility for downstream tasks. BPCloak [21] aligns RF features across behaviors, enabling identity authentication without exposing the underlying behavior. Conversely, ImmCOGNITO [24] reconstructs radar point clouds with a graph autoencoder, and AE-IPP [33] adds sample-specific adversarial perturbations to mmWave data, both enabling gesture or activity recognition without exposing user identity. MuFilter [23] exploits the invertibility of the FFT to selectively perturb predefined behavioral signatures in multi-user scenarios. P2ESA [16] uses an adversarially trained local encoder to retain information needed for authentication while obfuscating a predefined sensitive attribute. These methods target predefined privacy attributes and sensing tasks, but their outputs remain unencrypted. Communication eavesdroppers can therefore infer information that is not explicitly protected. Federated learning also reduces privacy leakage by avoiding centralized collection of raw sensing data [31, 53]. MIMOCrypt [26] uses MIMO encryption to protect multi-user WiFi sensing, but its communication-specific design cannot be directly applied to mmWave radar streams. SpecRing provides cryptographic access control over mmWave sensing semantics. Through reversible, key-derived spectrum shifts, range, Doppler, and angle can be encrypted independently, ensuring that each server can access only the dimensions authorized for its sensing task.

8 Conclusion

This paper presented SpecRing, a spectrum-shifting encryption framework for privacy-preserving mmWave sensing. SpecRing exploits the separable range, Doppler, and angle structures of FMCW signals and applies key-derived phase offsets to cyclically shift selected spectra. Authorized services recover the true physical scene through inverse rotation, while unauthorized parties observe a format-valid but semantically misleading virtual scene. To make this design practical, SpecRing further incorporated dimension-selective access control, environmental imprint removal, and pairwise key negotiation, preventing stable background structures and intercepted communications from exposing the secret offsets.

Our evaluation showed that SpecRing achieves exact recovery with the correct key and yields random-guessing-level performance with wrong keys. It preserves dimensional isolation, resists environmental imprint attacks, and incurs only low overhead. Overall, SpecRing demonstrated that spectrum shifting can serve as a lightweight and reversible primitive for semantic-level protection in mmWave sensing, enabling fine-grained privacy control while preserving the signal structure required by ubiquitous sensing services.

Acknowledgments

This work is supported in part by the National Natural Science Foundation of China under Grants 62602303, 62403349, 62372324, 62672257, in part by the Fundamental Research Funds for the Central Universities under Grants 63261080, 63251237, 63253266, 63263261, and in part by the Nankai University Education Foundation under Grant G02J0508.

References

[1] Suryoday Basak and Mahanth Gowda. 2022. mmSpy: Spying Phone Calls using mmWave Radars. In *Proc. IEEE S&P*. 1211–1228.

[2] Chandler Jackson Bauder, Tianhao Wu, Syed Irfan Ali Meerza, Aly Fathy, and Jian Liu. 2025. mm-RunAssist: mmWave-based Respiratory and Running Rhythm Analysis during Treadmill Workouts. *Proc. ACM IMWUT* 9, 3 (2025), 1–24.

[3] Mehdi Boukhechba, Elena Reynoso, Ioannis Pandis, Kenneth Mosca, Rumen Hristov, Shichao Yue, Yuqing Ai, Hariharan Rahul, Dina Katabi, Mark Morris, and Stefan Avey. 2023. Digital Phenotyping of Autoimmune Diseases Using Non-Contact Radio Frequency Sensing: A Longitudinal Study Comparing Systemic Lupus Erythematosus and Healthy Participants. In *Proc. ACM UbiComp/ISWC*. 615–619.

[4] Haoran Cao, Jiadi Yu, Hao Kong, Yi-Chao Chen, Yanmin Zhu, Linghe Kong, and Feilong Tang. 2025. OccMesh: Occlusion-aware Multi-user 3D Human Mesh Reconstruction Using mmWave Signals. *Proc. ACM IMWUT* 9, 3 (2025), 1–22.

[5] Qiming Cao, Hongfei Xue, Tianci Liu, Xingchen Wang, Haoyu Wang, Xincheng Zhang, and Lu Su. 2024. mmCLIP: Boosting mmWave-based Zero-shot HAR via Signal-Text Alignment. In *Proc. ACM SenSys*. 184–197.

[6] Jinbo Chen, Dongheng Zhang, Zhi Wu, Fang Zhou, Qibin Sun, and Yan Chen. 2024. Contactless Electrocardiogram Monitoring With Millimeter Wave Radar. *IEEE Transactions on Mobile Computing* 23, 1 (2024), 270–285.

[7] Lin Chen, Cong Li, Shuxin Zhong, Jun Chen, Yufei Wen, Haotian Song, and Kaishun Wu. 2026. Sensing Life in Stillness: Unified Dynamic and Static Human Mesh Reconstruction with mmWave Radar. *Proc. ACM IMWUT* 10, 1 (2026), 1–25.

[8] Xingyu Chen, Zhengxiong Li, Baicheng Chen, Yi Zhu, Chris Xiaoxuan Lu, Zhengyu Peng, Feng Lin, Wenyao Xu, Kui Ren, and Chunming Qiao. 2023. MetaWave: Attacking mmWave Sensing with Meta-material-enhanced Tags. In *Proc. ISOC NDSS*. 1–17.

[9] Tri Dao and Albert Gu. 2024. Transformers are SSMs: Generalized Models and Efficient Algorithms Through Structured State Space Duality. In *Proc. ICML*. 10041–10071.

[10] Pardis Emami-Naeini, Janarth Dheenadhayalan, Yuvraj Agarwal, and Lorrie Faith Cranor. 2023. Are Consumers Willing to Pay for Security and Privacy of IoT Devices?. In *Proc. USENIX Security*. 1505–1522.

[11] Guiyun Fan, Rong Ding, Xiaocheng Wang, Yichen Zhu, and Haiming Jin. 2025. m3ASL: ASL Gesture Recognition with Moving mmWave Radar. In *Proc. IEEE INFOCOM*. 1–10.

[12] Ruixu Geng, Jianyang Wang, Yuqin Yuan, Fengquan Zhan, Tianyu Zhang, Rui Zhang, Pengcheng Huang, Dongheng Zhang, Jinbo Chen, Yang Hu, and Yan Chen. 2026. A Survey of Wireless Sensing Security From a Role-Based View. *IEEE Communications Surveys & Tutorials* 28 (2026), 3412–3442.

[13] Alex Graves and Jürgen Schmidhuber. 2005. Framewise Phoneme Classification with Bidirectional LSTM and Other Neural Network Architectures. *Neural Networks* 18, 5–6 (2005), 602–610.

[14] Qingyong Hu, Qian Zhang, Hao Lu, Shun Wu, Yuxuan Zhou, Qianyi Huang, Huangxun Chen, Ying-Cong Chen, and Ni Zhao. 2024. Contactless Arterial Blood Pressure Waveform Monitoring with mmWave Radar. *Proc. ACM IMWUT* 8, 4 (2024), 1–29.

[15] David Hunt, Kristen Angell, Zhenzhou Qi, Tingjun Chen, and Miroslav Pajic. 2024. MadRadar: A Black-Box Physical Layer Attack Framework on mmWave Automotive FMCW Radars. In *Proc. ISOC NDSS*. 1–17.

[16] Andraž Krašovec, Gianmarco Baldini, and Veljko Pejović. 2025. P2ESA: Privacy-Preserving Environmental Sensor-Based Authentication. *Sensors* 25, 15 (2025), 4842.

[17] Yi Li, Chuyu Wang, Lei Xie, Qiancheng Jin, Long Fan, Jingyi Ning, and Sanglu Lu. 2024. Facial Landmark Detection Based on High Precision Spatial Sampling via Millimeter-wave Radar. *Proc. ACM IMWUT* 8, 4 (2024), 1–26.

[18] Zhengxiong Li, Fenglong Ma, Aditya Singh Rathore, Zhuolin Yang, Baicheng Chen, Lu Su, and Wenyao Xu. 2020. WaveSpy: Remote and Through-wall Screen Attack via mmWave Sensing. In *Proc. IEEE S&P*. 217–232.

[19] Guodong Liu, Lei Wang, Minjun Jiang, Qianran Qiao, Keran Li, Haipeng Dai, and Guihai Chen. 2026. LAM-assisted Acoustic Eavesdropping in Multi-speaker Scenarios via Commercial mmWave Radar. *Proc. ACM IMWUT* 10, 1 (2026), 1–28.

[20] Hankai Liu, Xiulong Liu, Xin Xie, Xinyu Tong, and Keqiu Li. 2023. PmTrack: Enabling Personalized mmWave-based Human Tracking. *Proc. ACM IMWUT* 7, 4 (2023), 1–30.

[21] Jianwei Liu, Chaowei Xiao, Kaiyan Cui, Jinsong Han, Xian Xu, and Kui Ren. 2023. Behavior Privacy Preserving in RF Sensing. *IEEE Transactions on Dependable and Secure Computing* 20, 1 (2023), 784–796.

[22] Xiulong Liu, Hankai Liu, Yantao Han, Xin Xie, Xinyu Tong, and Keqiu Li. 2025. MHTrack: mmWave-Based Mobile Hand Tracking. *IEEE Transactions on Mobile Computing* 24, 9 (2025), 9168–9183.

[23] Xiulong Liu, Hankai Liu, Jiaqi Zhang, Xin Xie, and Keqiu Li. 2025. Multi-User Behavioral Privacy Filtering for mmWave Radar Sensing. *IEEE Transactions on Mobile Computing* 24, 9 (2025), 8347–8361.

[24] Ying Liu, Si Zuo, Chao Yang, Yuqing Song, Dariush Salami, and Stephan Sigg. 2026. ImmCOGNITO: Identity Obfuscation in Millimeter-Wave Radar-Based Gesture Recognition for IoT Environments. *arXiv preprint arXiv:2602.07139* (2026).

[25] Ze Liu, Han Hu, Yutong Lin, Zhuliang Yao, Zhenda Xie, Yixuan Wei, Jia Ning, Yue Cao, Zheng Zhang, Li Dong, Furu Wei, and Baining Guo. 2022. Swin Transformer V2: Scaling Up Capacity and Resolution. In *Proc. IEEE/CVF CVPR*. 12009–12019.

[26] Jun Luo, Hangcheng Cao, Hongbo Jiang, Yanbing Yang, and Zhe Chen. 2024. MIMOCrypt: Multi-User Privacy-Preserving Wi-Fi Sensing via MIMO Encryption. In *Proc. IEEE S&P*. 2812–2830.

[27] Luoyu Mei, Ruofeng Liu, Zhimeng Yin, Qingchuan Zhao, Wenchao Jiang, Shuai Wang, Kangjie Lu, and Tian He. 2024. mmSpyVR: Exploiting mmWave Radar for Penetrating Obstacles to Uncover Privacy Vulnerability of Virtual Reality. *Proc. ACM IMWUT* 8, 4 (2024), 1–29.

[28] Chengzhen Meng, Chenming He, Dequan Wang, Yuxuan Xiao, Lingyu Wang, Xiaoran Fan, Lu Zhang, and Yanyong Zhang. 2025. GR-Fall: A Fall Detection System with Gait Recognition for Indoor Environments Using SISO mmWave Radar. *Proc. ACM IMWUT* 9, 3 (2025), 1–26.

[29] Rohith Reddy Vennam, Ish Kumar Jain, Kshitiz Bansal, Joshua Orozco, Puja Shukla, Aanjhan Ranganathan, and Dinesh Bharadia. 2023. mmSpoof: Resilient Spoofing of Automotive Millimeter-wave Radars using Reflect Array. In *Proc. IEEE S&P*. 1807–1821.

[30] Olaf Ronneberger, Philipp Fischer, and Thomas Brox. 2015. U-Net: Convolutional Networks for Biomedical Image Segmentation. In *Proc. MICCAI*. 234–241.

[31] Kanishka Roy, Tahsin Fuad Hasan, Chenfeng Wu, Eshwar Vangala, and Roshan Ayyalasomayajula. 2025. FedWiLoc: Federated Learning for Privacy-Preserving WiFi Indoor Localization. *arXiv preprint arXiv:2512.18207* (2025).

[32] Ville Satopaa, Jeannie Albrecht, David Irwin, and Barath Raghavan. 2011. Finding a "Kneedle" in a Haystack: Detecting Knee Points in System Behavior. In *Proc. IEEE ICDCSW*. 166–171.

[33] Biyun Sheng, Wangquan Qin, Hui Cai, Jun Li, Li Lu, Tie Qiu, and Fu Xiao. 2026. AE-IPP: Adversarial Example Enabled Identity Privacy Preserving with mmWave Signals. *IEEE Transactions on Mobile Computing* 25, 9 (2026), 15251–15266.

[34] Jayanth Shenoy, Zikun Liu, Bill Tao, Zachary Kabelac, and Deepak Vasisht. 2022. RF-Protect: Privacy against Device-Free Human Tracking. In *Proc. ACM SIGCOMM*. 588–600.

[35] Zhi Sun, Sarankumar Balakrishnan, Lu Su, Arupjyoti Bhuyan, Pu Wang, and Chunming Qiao. 2021. Who Is in Control? Practical Physical Layer Attack and Defense for mmWave-Based Sensing in Autonomous Vehicles. *IEEE Transactions on Information Forensics and Security* 16 (2021), 3199–3214.

[36] Mingyue Tang, Pranshu Teckchandani, Jizheng He, Hanbo Guo, and Elahe Soltanaghai. 2024. BSENSE: In-vehicle Child Detection and Vital Sign Monitoring with a Single mmWave Radar and Synthetic Reflectors. In *Proc. ACM SenSys*. 478–492.

[37] Jianyang Wang, Shiyuan Li, Di Cheng, Lingyun Zhou, Chang Chen, and Weidong Chen. 2023. CVAE: An Efficient and Flexible Approach for Sparse Aperture ISAR Imaging. *IEEE Geoscience and Remote Sensing Letters* 20 (2023), 1–5.

[38] Jianyang Wang, Dongheng Zhang, Binbin Zhang, Jinbo Chen, Yang Hu, and Yan Chen. 2024. RF-GymCare: Introducing Respiratory Prior for RF Sensing in Gym Environments. *Proc. ACM IMWUT* 8, 3 (2024), 1–28.

[39] Liang Wang, Aoyang Chang, Boqun Liu, Zhiwen Yu, Pengfei Hu, Jia Liu, Yao Zhang, and Bin Guo. 2025. mmHSE: Enhanced Eavesdropping Attack on Headsets Leveraging COTS mmWave Radar. *Proc. ACM IMWUT* 9, 2 (2025), 1–27.

[40] Lingyu Wang, Kai Wang, Dequan Wang, You Zuo, Chenming He, Chengzhen Meng, Xiaoran Fan, Haojie Ren, and Yanyong Zhang. 2025. mmMUSE: An mmWave-based Motion-resilient Universal Speech Enhancement System. *Proc. ACM IMWUT* 9, 4 (2025), 1–33.

[41] Maximiliane Windl, Omer Akgul, Nathan Malkin, and Lorrie Faith Cranor. 2025. Privacy Solution or Menace? Investigating Perceptions of Radio-Frequency Sensing. In *Proc. USENIX Security*. 6045–6064.

[42] Weiye Xu, Danli Wen, Jianwei Liu, Zixin Lin, Yuanqing Zheng, Xian Xu, and Jinsong Han. 2025. DiskSpy: Exploring a Long-Range Covert-Channel Attack via mmWave Sensing of μm-level HDD Vibrations. In *Proc. USENIX Security*. 3903–3920.

[43] Xiangyu Xu, Yu Chen, Zhen Ling, Li Lu, Junzhou Luo, and Xinwen Fu. 2024. mmEar: Push the Limit of COTS mmWave Eavesdropping on Headphones. In *Proc. IEEE INFOCOM*. 351–360.

[44] Meng Xue, Wentao Xie, Zuohuizi Yi, Zhilong Zhang, Shumao Wu, Yinan Zhu, Qian Zhang, and Changzheng Chen. 2025. Home-based Dry Eye Assessment via Blink Kinematics Using mmWave and Clinical Knowledge Distillation. In *Proc. ACM MobiCom*. 1121–1135.

[45] Hongliu Yang, Duo Zhang, Xusheng Zhang, Jie Xiong, Zizhou Fan, Wanru Ning, Weiyan Chen, Fusang Zhang, Zijun Han, and Daqing Zhang. 2025. From Spatial Domain to Temporal Domain: Unleashing the Capability of CFAR for mmWave Point Cloud Generation. *Proc. ACM IMWUT* 9, 2 (2025), 1–29.

[46] Junlin Yang, Jiadi Yu, Linghe Kong, Yanmin Zhu, and Hong-Ning Dai. 2025. Toward Open-World-Aware User Authentication Based on Human Bodies Using mmWave Signals. *IEEE Transactions on Mobile Computing* 24, 9 (2025), 9036–9049.

[47] Muhammet Emin Yanik and Murat Torlak. 2019. Near-Field 2-D SAR Imaging by Millimeter-Wave Radar for Concealed Item Detection. In *Proc. IEEE RWS*. 1–4.

[48] Yuqin Yuan, Jinbo Chen, Dongheng Zhang, Ruixu Geng, Hanqin Gong, Guixin Xu, Yu Pu, Zhi Lu, Yang Hu, Dong Zhang, Likun Ma, Qibin Sun, and Yan Chen. 2025. Atrial Fibrillation Detection via Contactless Radio Monitoring and Knowledge Transfer. *Nature Communications* 16 (2025), 4317.

[49] Duo Zhang, Xusheng Zhang, Zhehui Yin, Pengfei Zhou, Hongliu Yang, Pei Wang, Zhiyun Yao, Junzhe Wang, Fusang Zhang, and Daqing Zhang. 2025. Breaking the Resolution Barriers of mmWave Arrays via Null Steering for Sleep Monitoring in Multi-Person Scenarios. *Proc. ACM IMWUT* 9, 3 (2025), 1–29.

[50] Fenglin Zhang, Zhebin Zhang, Le Kang, Anfu Zhou, and Huadong Ma. 2025. mmTAA: A Contact-Less Thoracoabdominal Asynchrony Measurement System Based on mmWave Sensing. *IEEE Transactions on Mobile Computing* 24, 2 (2025), 627–641.

[51] Rui Zhang, Ruixu Geng, Yadong Li, Ruiyuan Song, Hanqin Gong, Dongheng Zhang, Yang Hu, and Yan Chen. 2025. RFMamba: Frequency-Aware State Space Model for RF-Based Human-Centric Perception. In *Proc. ICLR*. 1–17.

[52] Shichen Zhang, Qijun Wang, Maolin Gan, Zhichao Cao, and Huacheng Zeng. 2025. RadSee: See Your Handwriting Through Walls Using FMCW Radar. In *Proc. ISOC NDSS*. 1–17.

[53] Cui Zhao, Guotong Fang, Han Ding, Xinhui Liu, Fei Wang, Ge Wang, Kun Zhao, Zhi Wang, and Wei Xi. 2025. Federated Multi-Source Domain Adaptation for mmWave-Based Human Activity Recognition. *IEEE Transactions on Mobile Computing* 24, 8 (2025), 7283–7296.

[54] Running Zhao, Luca Jiang-Tao Yu, Tingle Li, Zhihan Jiang, Chenwei Zhang, Chenshu Wu, Hang Zhao, and Edith C.H. Ngai. 2026. SPACE: Speaker Adaptation for Acoustic Eavesdropping Using mmWave Radio Signals. *IEEE Transactions on Mobile Computing* 25, 1 (2026), 723–738.

[55] Rui Zhou, Songlin Li, Hongwang Zhang, Chenxu Liu, and Jiajun Sun. 2025. mmMulti: Multi-person Action Recognition Based on Multi-task Learning Using Millimeter Waves. *Proc. ACM IMWUT* 9, 2 (2025), 1–25.