NeuroRadar: A Neuromorphic Radar Sensor for Low-Power IoT Systems

Kai Zheng
University of California San Diego
kazheng@ucsd.edu

Kun Qian
University of California San Diego
kuq002@ucsd.edu

Timothy Woodford
University of California San Diego
twoodfor@ucsd.edu

Xinyu Zhang
University of California San Diego
xyzhang@ucsd.edu

ABSTRACT

Radar sensors have recently been explored in the industrial and consumer Internet of Things (IoT). However, such applications often require self-sustainable or untethered operations, which are at odds with the high power consumption of radar. This paper proposes NEURORADAR, a neuromorphic radar sensor, to achieve low-power wireless sensing. NEURORADAR jointly optimizes the analog hardware and the computation model, in order to mimic the highly efficient biological sensing and neural processing system. NEURORADAR features a highly simplified radar front end, which eliminates the power-hungry components in conventional radars. It directly "encodes" ambient motion into spiking signals, which can be processed using spiking neural networks running on energy-efficient neuromorphic computing platforms. We have prototyped NEURORADAR and evaluated its performance in two use cases: gesture sensing and localization. Our experiments demonstrate that NEURORADAR can achieve high sensing accuracy, at orders of magnitude lower power consumption compared with traditional radar.

CCS CONCEPTS

• **Human-centered computing** → **Ambient intelligence**; **Ubiquitous computing**.

KEYWORDS

Neuromorphic computing, Neuromorphic sensors, Spiking neural networks, Low-power sensing, Gesture recognition, Motion tracking.

ACM Reference Format:
Kai Zheng, Kun Qian, Timothy Woodford, and Xinyu Zhang. 2023. NeuroRadar: A Neuromorphic Radar Sensor for Low-Power IoT Systems. In *ACM Conference on Embedded Networked Sensor Systems (SenSys '23), November 12–17, 2023, Istanbul, Turkiye.* ACM, New York, NY, USA, 14 pages. https://doi.org/10.1145/3625687.3625788

1 INTRODUCTION

Radar sensors in the Internet of Things (IoT) systems have gained traction in recent years, and are widely used in healthcare, smart homes, industrial automation, and intelligent transportation [28, 55, 58]. By 2025, industrial radar applications are anticipated to encompass 10 million devices, whereas the consumer market will reach a substantial $250 million [93]. Nevertheless, the high power consumption of radar hardware remains a significant challenge, particularly for battery-operated IoT devices and wearables where energy efficiency and battery lifespan are crucial. Compounding this issue, numerous smart sensing applications, such as motion-activated security radar, wearable gesture recognition, and activity classification, often employ power-intensive artificial neural networks (ANNs) for signal processing. Unlike human neurons that operate in short, pulse-based bursts, ANNs prolong the activity of their "neurons" using continuous activation functions, which substantially increases the power demands of IoT devices. Furthermore, ANNs utilize the classical von Neumann architecture, which frequently shuttles data between physically separate CPU and memory units, resulting in additional processing overhead.

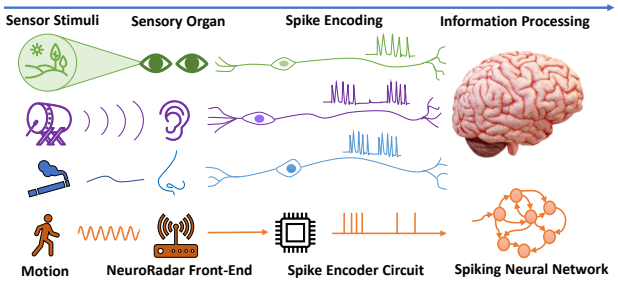

Figure 1: Analogy of neuromorphic radar sensor. NEURORADAR achieves energy-efficient sensing by emulating the structure and functionality of biological sensing systems.

Recent advancements in neuromorphic engineering have inspired Spiking Neural Networks (SNNs) [47] and dedicated neuromorphic circuits [18] that better approach the efficiency of sensory signal processing in the brain. SNNs are structured to mirror the pulse-based behavior of the human nervous system. They consist of spiking neurons and the synaptic connections between them. Realized on dedicated neuromorphic circuits, SNNs showcase exceptional energy efficiency that surpasses traditional von Neumann computing units by orders of magnitude [10]. The revolution in neuromorphic computing has also given rise to state-of-the-art neuromorphic sensing hardware, such as the energy-efficient, fast-response event camera [40].

Inspired by these advancements, recent research has proposed SNN-based signal processing to facilitate low-power radar operation

SenSys '23, November 12–17, 2023, Istanbul, Turkiye

ACM ISBN 979-8-4007-0414-7/23/11.
https://doi.org/10.1145/3625687.3625788

[6, 7]. However, these systems do not incorporate a full-fledged neuromorphic hardware architecture. Primarily, the analog front-end of these "SNN radar" systems [6, 7] remains the same as traditional radars. Although the SNN-based signal processing has lowered the signal processing power consumption to the order of hundreds of μW [76], the radar front-end can demand tens to hundreds of mW. This discrepancy poses a challenge to achieving truly energy-efficient radar sensing. Additionally, the "SNN radar" systems [6, 7] continue to rely on conventional CPUs or digital signal processing (DSP) units for signal processing. The radar signals have to be first sampled by analog-to-digital converters (ADCs), mapped into spikes, and then processed by SNNs for ranging or environmental perception. Unfortunately, the extra sampling steps prior to the SNN involve traditional computing units, which adds a substantial overhead, underutilizing neuromorphic computing's full potential.

In this paper, we introduce NEURORADAR, a novel low-power radar sensing system that fully exploits the power of neuromorphic sensing and computing. NEURORADAR draws inspiration from neuromorphic sensors that mimic mammalian sensory systems, generating event-triggered outputs in response to external stimuli, as depicted in Fig. 1. Contrary to traditional radars with continuous frame-based outputs, NEURORADAR produces spiking patterns upon detecting motion in the surrounding area. Unlike the recently proposed SNN radars [6, 7], NEURORADAR follows a neuromorphic architecture that jointly designs the analog sensing front-end and spiking signal processing:

(1) SIL-based radar sensor front-end. NEURORADAR employs a drastically simplified RF front-end that removes most power-intensive active RF components that exist in traditional radars, leaving only a low-power free-running oscillator. NEURORADAR senses environmental changes using the self-injection locking (SIL) principle [90], where the oscillator's frequency is influenced by motion in the surrounding area. Despite the simplistic design, NEURORADAR preserves a reasonable level of sensitivity due to the inherent properties of SIL architecture that enhance the signal strength. However, a single SIL sensor is unable to provide angular resolution and accurate range information, as it only senses environmental motion information. To overcome this limitation, we draw inspiration from the compound structures found in certain biological eyes [31], and propose to design an array of SIL sensors with judiciously separated carrier frequencies. With the SIL sensor array, NEURORADAR can implicitly encode spatial information through the multi-channel spiking signals, which can subsequently be decoded using application-specific SNN models.

(2) Analog spike encoding and full SNN processing. NEURORADAR converts ambient motion signals from the sensor front-end into spikes using an analog spike encoding circuit. The spike encoder follows a biological neuron model, preserving all the essential sensing information in the spike sequences. The spike sequences can then be directly processed by the SNNs on neuromorphic computing systems, thereby eliminating the need for any non-spike-based computing units. Consequently, we can train the SNNs using these raw spike signals for various tasks, including gesture recognition and localization. This comprehensive SNN processing workflow allows NEURORADAR to deliver application-specific sensing results with superior energy efficiency.

To verify the effectiveness of our design, we prototype NEURORADAR using discrete RF circuits and further perform simulation for the integrated circuit (IC) version. Our experiments show that a single-RF-chain NEURORADAR can effectively sense motion in the environment while consuming only 780 μW power (IC: 240 μW), which is 1-2 orders of magnitude lower than existing continuous wave radar systems with similar operating frequencies. We further conduct two case studies to verify the usability of NEURORADAR for practical IoT sensing applications. Specifically, NEURORADAR can facilitate hand gesture recognition with an accuracy of 94.6% and perform moving target localization with an average error of 0.98 m. Compared with other SNN-based gesture recognition systems [6, 7, 67] with similar capability, NEURORADAR saves 78%-93% computing power. NEURORADAR reduces the end-to-end power consumption by 1-2 orders of magnitude for both use cases, compared with existing radars. Considering the spatial resolution and motion detection capabilities, NEURORADAR can potentially be used in a wide range of wireless IoT sensing applications, such as vital sign sensing, surveillance alarm, *etc.*

In summary, we make the following contributions:

(i) We introduce NEURORADAR, a novel low-power radar paradigm that realizes the concept of neuromorphic radar sensing. NEURORADAR incorporates a spike-generation radar sensor that directly interfaces with SNN-based neuromorphic processors, leading to superior energy efficiency.

(ii) We devise a low-power, low-complexity radar front end based on the SIL principle. Both our theoretical analysis and experimental results demonstrate that multi-chain SIL radar sensors can supply ample information for short-range, low-velocity sensing applications.

(iii) We implement the neuromorphic radar system through a PCB prototype and carry out simulations for the IC version. Our experiments verify NEURORADAR's capability to empower the resource-constrained IoT devices to perform low-power smart sensing.

2 BACKGROUND

2.1 Self-Injection Locking

Self-injection locking [14] is a phenomenon where an oscillator's frequency is affected by a reflected version of its own signal, as depicted in Fig. 2. Unlike conventional injection locking, where the oscillator's frequency is locked to the frequency of an external injection signal [1], *the frequency of a self-injection-locked oscillator (SILO) is dependent on the amplitude and phase of the reflected signal*.

Based on classical analysis of injection locking [60], we can model the frequency shift of the oscillator caused by the reflectors in the environment:

$$\Delta\omega \approx \sum_{n=1}^{N} \frac{\omega_o}{2Q} \frac{|I_{inj}^{(n)}|}{|I_{osc}|} sin(\omega_o \cdot \frac{2r^{(n)}}{c}) \quad (1)$$

Here, ω_o is the center frequency of the oscillator; Q is the quality factor of the LC resonating tank; $I_{inj}^{(n)}$ is the injection signal from the n^{th} reflector; I_{osc} is the oscillator signal; $r^{(n)}$ is the distance of

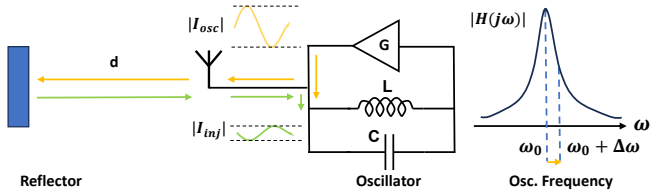

Figure 2: Self-injection locking.

the n^{th} reflector; c is the speed of light, and N is the total number of reflectors.

Notably, the strength of the reflected signal, $|I_{inj}|$, is proportional to $1/r^2$. Its phase, $\angle I_{inj}$, encapsulated in the sine term, is also related to r. *Thus, the oscillation frequency is modulated by the motion of the reflectors*. In the case of *multiple reflectors*, the observed frequency shift of a SILO is the summation of the shifts caused by each reflector. Since static reflectors cause constant frequency shifts, *only moving reflectors contribute to frequency modulation*. This principle is leveraged by NEURORADAR for environmental perception.

2.2 Spiking Neural Networks

Biological neurons communicate by generating and propagating electrical pulses or spikes. Neurons are interconnected via specialized junctions termed synapses. A neuron fires a *spike* whenever enough incoming pulses accumulate to push its membrane potential above a certain threshold, following which the neuron resets itself. This process is often abstracted as Leaky-Integrate-and-Fire (LIF) [26]. In traditional ANNs, neurons encode information in a complex network of real-valued activations. Activation functions such as ReLU essentially approximate the spiking rates of biological neurons. In contrast, SNNs mimic the human neuron system more closely by: (i) using spiking signals directly for inter-neuronal communication and (ii) using the *timing rather than shape* of spikes to convey neural information.

The computation and energy efficiency advantages of SNNs originate from two fundamental aspects. First, the neuromorphic architecture can realize massive parallel processing, since each neuron represents an integrated memory and computation unit, in contrast to the rigid separation of CPU and RAM in von Neumann architectures. Thus, SNNs can potentially continue to push the "intelligence per Joule" as Moore's law scaling comes to an imminent end. Secondly, the energy consumption of SNNs is proportional to the number of processed spikes, with each spike requiring as little as a few pico Joules [10]. As information is sparsely encoded in the rates/timing of the spiking neurons, an SNN can implement the same end-to-end functionality as an ANN [69] but with much lower energy expenditure.

Notably, the advantages of SNNs can be manifested only on specialized non-von Neumann in-memory computing platforms specifically designed to process spiking inputs. The past decade has witnessed a variety of such platforms, such as Intel Loihi [18], uBrain [76], DynapSE [53], IBM Truenorth [3], and SPINNaker [24]. Albeit an active area of research, neuromorphic computers have already demonstrated orders of magnitude of energy efficiency than conventional computing architectures [10].

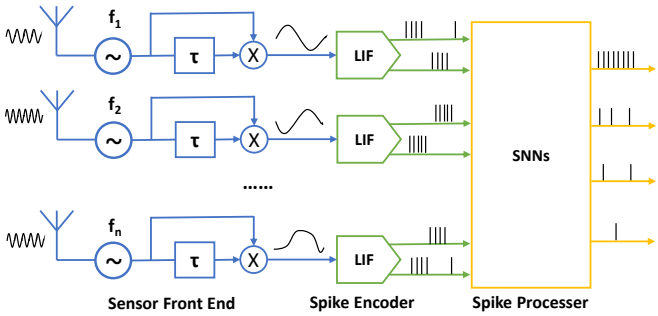

Figure 3: System Overview.

3 SYSTEM OVERVIEW

NEURORADAR consists of three main components: sensor front-end, spike encoders, and spike processors (Fig. 3). The sensor front-end senses ambient motion and the output signals are converted into spike sequences (referred to as *spike trains*) by the spike encoders. These spike trains are then directly processed by the energy-efficient SNNs.

Sensor front-end. The NEURORADAR front-end emits a weak, continuous-wave single-tone signal in the 0.3~3 GHz ultra-high frequency (UHF) band. The core component is a SILO whose frequency is modulated by the motion of the surrounding targets [90]. By demodulating this frequency shift, the system generates a baseband signal that carries the motion information. We further introduce a sensor array design that combines multiple SILOs with different operating frequencies to provide richer spatiotemporal information.

Spike encoder. The spike encoding circuit takes the baseband signal produced by the front-end and converts the signal into spike trains following the LIF model [26] (Sec. 2.2). Given that the input is AC-coupled and the signal comprises both positive and negative parts, two spike encoders are jointly employed to encode each channel of the radar sensor. The spike encoding circuits operate entirely in an *event-driven* manner; they only generate spikes when the sensor front-end detects motion and stays idle otherwise.

Spike processor. The spike encoders interface directly with the neuromorphic computing circuits, enabling all signals to be processed within the spike domain. Our approach involves designing multi-layer convolutional SNNs to process the multi-channel spike chains from the NEURORADAR sensor array. These SNNs execute pattern recognition and regression tasks according to the application requirements.

4 SENSOR FRONT-END DESIGN

4.1 Design Principle

The main principle of the front-end design is to reduce its power consumption for NEURORADAR. To achieve it, we first analyze the power-hungry RF components of traditional radars that lead to high power consumption. A typical continuous-wave (CW) radar front-end, as shown in Fig. 4(a), includes elements such as voltage-controlled oscillator (VCO), phase-locked-loop (PLL), crystal oscillator (XO), mixer, low-noise amplifier (LNA), and power amplifier (PA). While power consumption can vary depending on specific designs, we annotate a representative CW radar [86] for reference. These active RF components are necessary to maintain the high

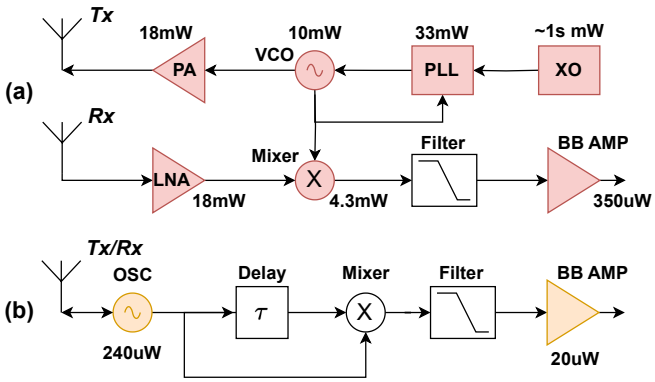

Figure 4: Radar front-end architecture comparison and power consumption breakdown. Uncolored blocks are passive components. (a) A typical CW radar [86]. (b) NEURORADAR.

sensing performance required for advanced applications such as automotive perception. For instance, the phase noise of a radar directly impacts target detectability, spatial resolution, and maximum range [20, 72]. To contain the phase noise, most radar systems use a VCO and PLL in a feedback loop, using a high-precision XO as the reference input to synthesize a low-phase noise radar signal with a precise frequency. Consequently, such high-profile radar front-ends require a high power budget of several hundred mW, irrespective of the signal processing hardware.

In contrast to traditional radar systems, neuromorphic systems exhibit superior power efficiency and rapid response times by emulating the *event-driven* communication and computation in biological neural systems [11, 44]. Event cameras [40], also known as dynamic vision sensors, represent an epitome of neuromorphic sensing systems. Instead of capturing full frames at a fixed rate, event cameras generate asynchronous events in response to changes in pixel-level brightness. This event-driven approach increases the camera's dynamic range, while substantially reducing power consumption and data processing load [40].

Inspired by the event camera, we design NEURORADAR sensor front-end that only responds to *changes* in the radar channel (caused by motion) and produces asynchronous spike signals that contain relevant motion information. To attain these properties, we extend the SIL structure and develop spike encoders to convert radar signals into spike trains. A SIL radar, as a variant of Doppler radar [88], is inherently a motion detector, which aligns well with the event-driven neuromorphic sensing principle. Moreover, radars with more complex waveforms, like wide-band Frequency-Modulated Continuous Wave (FMCW) radar, require intricate baseband signal processing (such as FFTs) to extract basic sensing information. This is challenging to implement without ADC sampling and conventional DSP units. In contrast, a SIL radar emits single-tone signals and "demodulates" motion directly from the reflected signals, which eliminates the need for complicated wideband signal processing and facilitates the use of spike encoders. SIL radars feature a simplified architecture, which makes them power-efficient and cost-effective to implement. We elaborate on the SIL front-end design of NEURORADAR in the following sections.

4.2 SIL Sensor Design

SIL radar adopts a simplistic architecture with only three RF components: an oscillator, a time delay unit, and a mixer (Fig. 4(b)). The oscillator emits an RF signal that becomes self-injection-locked due to environmental reflections. A time delay unit and mixer demodulate the frequency shift caused by moving targets. The system's total power consumption is kept under $300\mu W$ by lowering the oscillator's output power and employing a passive low-power mixer.

While the removal of active RF components such as LNA or PA typically results in low sensitivity, the SIL radar's unique architecture provides a sensitivity gain that compensates for the impact. This property ensures that despite the simplified design, a SIL radar can still support our targeted IoT sensing applications that require only limited range/velocity resolution (*e.g.*, occupancy detection, coarse indoor tracking, hand gesture recognition).

A target's motion induces *phase modulation* on conventional a Doppler radar [56], in contrast to *frequency modulation* on the SIL radar [90]. The demodulation circuit extracts the phase change over the delay time τ_d. As phase is the time integral of frequency, SIL radar's demodulation process inherently integrates and enhances the motion signal [79]. A larger *demodulation gain* and hence high *sensitivity* can be achieved by increasing τ_d, if the motion frequency is much lower than $1/\tau_d$, which holds true for our targeted IoT use cases. Furthermore, in SIL radar, the oscillator signal, which contains sensing information, directly enters the mixer without any attenuation. In contrast, for conventional Doppler radar, the mixer's input comes from the attenuated reflected signal. Therefore, *SIL radar effectively amplifies the received signal amplitude to the oscillator output level for free*. Following the empirical model in [79], we find that the SIL radar can provide a sensitivity gain of around 19.97 dB with $\tau_d = 80$ ns and carrier frequency 915 MHz (corresponding to our implementation), which can be traded for low-power operations.

4.3 Array of SIL Sensors

4.3.1 Sensing Information from a Single SIL Radar. Suppose a target is moving randomly within the surrounding area and a continuous frequency shift (Eq. (1)) of the oscillator is observed. The demodulation circuit integrates the frequency shift $\Delta\omega(t)$ during the delay time τ_d and the demodulated output $y(t) \approx \Delta\omega(t)\tau_d$, if $\Delta\omega(t)$ changes slowly and is considered constant over τ_d. Therefore $y(t)$ is a direct representation of the motion-modulated frequency signal and from Eq. (1),

$$y(t) = a(t)sin[4\pi f\frac{r(t)}{c}] \tag{2}$$

$$a(t) = G_d\frac{\omega_0}{2Q}\frac{1}{|I_{osc}|}\frac{\sigma}{r^2(t)}, \tag{3}$$

Here G_d represents an abstracted gain encapsulating various factors (*i.e.* demodulation gain, antenna gain, and any practical system loss), and σ is the radar cross section (RCS) of the target, which is an unknown parameter and may fluctuate over time.

The range information $r(t)$ is embedded in both $a(t)$ and $sin[4\pi f\frac{r(t)}{c}]$, but this alone is insufficient for localizing the target. Although $a(t)$ includes absolute range information, $r(t)$ cannot be estimated due to the unknown target RCS σ. Additionally, as $sin[4\pi f\frac{r(t)}{c}]$ is

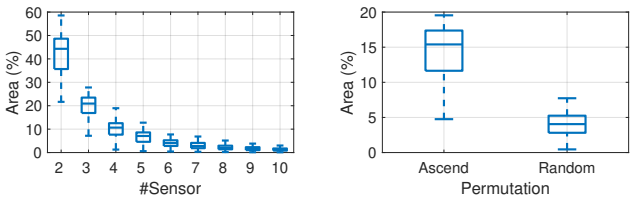

(a) Different number of sensors (b) Different frequency permutation

Figure 5: Localization ambiguity area with different settings.

$2\pi-$periodic, it only contains *ambiguous range information*. Moreover, a single sensor *fails to provide the angular information* of the target. This implies that given a distance r, the actual location of the target could be anywhere on a circle with a radius of r. Therefore, further information is required to precisely localize the target.

4.3.2 Frequency-Diverse SIL Sensor Array Design. To overcome the lack of range/angle resolution, NEURORADAR combines multiple SIL radar sensors operating at different frequencies, forming a frequency-diverse array (FDA). A colocated sensor array can infer the direction of the target by exploiting the phase difference of the received signals across the sensors, whereas the frequency diversity offers the potential to resolve the range ambiguity. Distinct from traditional antenna arrays that often employ a half-wavelength ($\lambda/2$) spacing to avoid angular aliasing [78], NEURORADAR employs a quarter-wavelength ($\lambda/4$) spacing, λ being the average wavelength, because a SIL radar uses the same antenna for both transmission and reception, effectively doubling the phase difference due to antenna spacing.

To quantize the sensing capability of the array design, we derive a model-driven localization and speed estimation process. Consider a linear array of K sensors, where the position and frequency of the k-th sensor are $\vec{d}_k$ and f_k, respectively. Since NEURORADAR only detects motion, we suppose a target moves at a constant speed within a short time (*e.g.*, 0.5 s), and a total of M observations are made with an interval of Δt. According to Eq. (2), when a target is located at $\vec{l}$ and moves with velocity $\vec{v}$, the theoretical observation vector:

$$\vec{s}(\vec{l},\vec{v}) = [\sin(4\pi f_k \frac{r_{k,m}}{c})]_{k=0,m=0}^{k=K-1,m=M-1}, \tag{4}$$

where $r_{k,m} = ||\vec{l}+\vec{v}\cdot m\Delta t - \vec{d}_k||_2$ is the distance between the target's location at time $m\Delta t$ and the k-th radar sensor at $\vec{d}_k$. In the model, the amplitude $a(t)$ is not considered because in practice $a(t)$ fluctuates randomly due to the time-varying RCS and provides unreliable information.

Given the real observation vector $\vec{y} = [y_{k,m}]$ from the radar sensor array, the location and speed of the target can be estimated by

$$\vec{l}_{\text{opt}}, \vec{v}_{\text{opt}} = \text{argmax}_{\vec{l},\vec{v}} \frac{\vec{s}^T(\vec{l},\vec{v})\cdot\vec{y}}{|\vec{s}(\vec{l},\vec{v})|\cdot|\vec{y}|}. \tag{5}$$

Here $[\vec{s}^T(\vec{l},\vec{v})\cdot\vec{y}]$ computes the correlation between the real observation vector and the theoretical observation vector.

To establish the *optimal number of radar sensors* required to achieve reasonable resolution, we use the above model to numerically derive the sensing resolution. In the numerical simulation, the sensor operating frequencies are set between 800 MHz and 950 MHz.

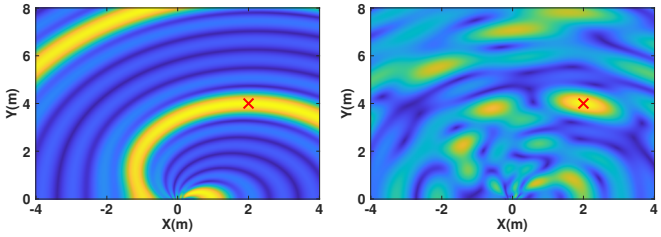

(a) Ascending permutation (b) Random permutation

Figure 6: Localization ambiguity area. The true target location is marked with ×.

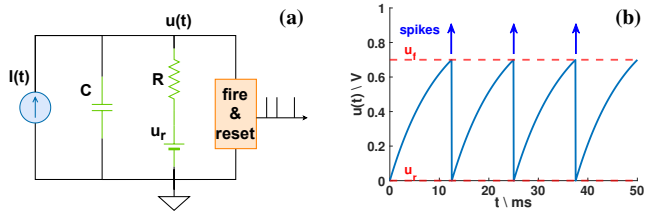

Figure 7: (a) RC circuit model for a LIF neuron. (b) The evolution of the membrane potential with a constant current injection.

The simulation considers targets ranging from 0.5 to 7 meters and sensing directions between 45° and 135° relative to the sensor array. Target speeds range from 0.5 to 3 m/s, with their moving directions varying between 0° and 360°. We randomly sample 100 targets and calculate the average percentage of the area with a correlation value (Eq. (5)) exceeding -3 dB. This area represents the *ambiguity* of NEURORADAR.

The results depicted in Fig. 5a reveal that the *ambiguous area decreases as the number of sensors increases*. A smaller ambiguous area signifies a reduction in ambiguous side lobes and a more concentrated main lobe. The empirical findings suggest that *an array of 6 sensors is sufficient for* NEURORADAR *to resolve most targets*, striking an effective balance between resolution and array size.

The configuration of frequencies in the sensor array impacts sensing ambiguity. As shown in Fig. 5b, a random frequency permutation results in a significantly smaller ambiguous area compared to an ascending permutation. This finding aligns with previous research on traditional FDA [2, 48], which indicated that nonlinear and random frequency offsets result in a range-angle decoupled beam pattern. Fig. 6 presents the location ambiguity area given an observation vector from a target at a specific location. In this instance, we only consider the location for simplicity, assuming that the observation vector can be retrieved even if the target is stationary. Consequently, NEURORADAR takes advantage of this property by *adopting a nonlinear frequency offset and a random frequency permutation* for its sensor array.

5 SPIKE ENCODING AND PROCESSING

5.1 Spike Encoder Design

To facilitate end-to-end SNN signal processing, NEURORADAR employs an analog spike encoding circuit to directly transform the SIL radar signals into spike trains. The spike encoder must preserve the essential sensing information in Eq. (2). To this end, the encoder should perform *spike rate encoding*, in which the firing frequency

increases linearly with the amplitude of the input signal. As a result, *the phase information in Eq.* (2) *is represented as variations in spike density*.

We design our spike encoder based on the aforementioned LIF neuron model [26] (Sec. 2.2). The LIF neuron model consists of a current injector, an RC parallel circuit, and a spike firing circuit, as depicted in Fig. 7a. In the human nervous system, a neuron's membrane potential $u(t)$ rises upon receiving input stimuli $I(t)$ from other neurons. Once $u(t)$ reaches a threshold u_f, the neuron triggers a spike to adjacent neurons and resets its voltage to a resting value u_r, as shown in Fig. 7b. In the absence of input, the membrane potential decays exponentially to its resting value through a leaky resistance path.

The evolution of $u(t)$ can be characterized as:

$$\tau_m \frac{du(t)}{dt} = -[u(t) - u_r] + RI(t) \tag{6}$$

$$\lim_{\delta \to 0; \delta > 0} u(t+\delta) = u_r, \text{ when } u(t) = u_f \tag{7}$$

Here, the membrane time constant $\tau_m = RC$ determines the decay time of the membrane voltage, with C being the membrane capacitance and R representing the leaky resistance.

Given a constant input I_0 from the SIL radar, the spike firing interval can be found by solving the differential Eq. (6),

$$T_s = -\tau_m \, log(1 - \frac{u_f - u_r}{RI_0}) \tag{8}$$

When the leaky resistor is large, the spike firing rate $f_s = \frac{1}{T_s} \approx \frac{I_0}{(u_f - u_r)C}$, which grows linearly with the input signal. The leaky resistor causes a small input *dead zone* where no spike is fired even the input $I_0 > 0$, which is implied by Eq. (8) that $I_0 > \frac{u_f - u_r}{R}$ must be satisfied. We utilize this property and design a proper dead zone to suppress random noise input and avoid spike misfiring.

To encode a continuous signal into a spike sequence without information loss, a sufficiently large spike rate is required. Given a spike sequence with spike time t_i, and signal bandwidth Ω,

$$t_{i+1} - t_i < \frac{\pi}{\Omega}, \forall \, i \tag{9}$$

must be satisfied to guarantee perfect recovery [15]. As a larger spike firing rate increases power consumption, our design strikes a tradeoff between input bandwidth and power consumption.

5.2 SNN Design

As the motion signals are converted into multiple parallel spike trains, we design SNNs to process the spiking signals and extract the spatiotemporal features. The overall structure of the SNN includes three main components: spike buffering units, convolution layers, and spike decoders (Fig. 8).

Spike buffering units. The input spike sequences initially arrive at the spike buffering units, which are made up of cascaded time delay units. Each delay unit imposes a consistent time delay of n_{dly} clock ticks, and the output spikes then enter the next-stage time delay unit. In the majority of neuromorphic computing hardware, SNNs are realized using digital circuits, with neuron states being updated synchronously according to a clock tick (*e.g.*, 1 ms). Upon the completion of the input sequence, the spike buffering units concatenate the outputs from all delay units and present the spikes concurrently

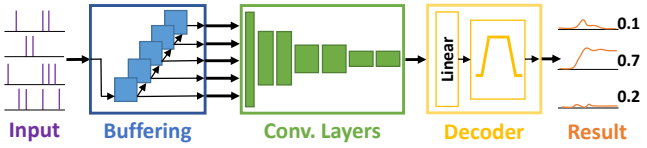

Figure 8: Illustration of the SNN structure.

to the subsequent layer. To improve the performance of the SNN, the buffered spikes are repetitively dispatched to the next layer every n_{int} clock ticks. By flattening the temporal dimension of the spike sequence, the spike buffering units simplify the task for the subsequent convolution layers in extracting the temporal features of the spike sequence. A similar method is employed in [4] for processing spike data from the event cameras.

Convolution layers. Convolutional layers are essential components of Convolutional Neural Networks (CNNs) which can detect local spatial patterns and structures within an image. With the spike buffering units flattening the temporal dimension of the input spike sequences, convolution layers can be similarly employed to extract the spatiotemporal features of the spike sequences. Consequently, we design a stack of convolution layers, accompanied by other types of layers such as pooling layers and fully connected layers, to process and classify the extracted features.

Spike decoders. In NeuroRadar, the SNNs are trained in such a way that the output values are represented by the spike firing rate of neurons in the final layer. Eventually, the output spike rate has to be converted into a continuous value that can be interpreted by the sensing applications. For classification tasks, the prediction probability for each class can be determined by applying low-pass filtering to the spikes from each output neuron representing the respective classes. For regression tasks, the output values are represented by an ensemble of neurons. We train decoders to perform a linear mapping between the neuron outputs and the final output following [75].

5.3 SNN Training

SNN training is crucial for extracting spatiotemporal features from input data. In NeuroRadar, the trainable parts of the SNNs are the convolution layers and the spike decoders. The ANN-SNN conversion method [65] is employed for training the SNNs in NeuroRadar. The method involves training a conventional ANN with the same structure as the desired SNN. Given an initial spike sequence, some spike buffering units fire spikes into the next layer at a constant firing rate, while others do not fire spikes at all. This allows the conversion of the input into a static image with 0/1 binary pixels. We then utilize conventional neuron models, like ReLU neurons, and conventional backpropagation algorithms to optimize the connection weights within the ANN. After training, all the ReLU neurons in the ANN are replaced with spiking neuron models, specifically LIF neurons. Lastly, weight scaling needs to be performed for the SNN to ensure a reasonable spike firing frequency. After completing these steps, the trained ANN is effectively converted into an SNN, which can then process spike input efficiently and accurately.

6 SYSTEM IMPLEMENTATION

We build a NeuroRadar hardware prototype using discrete components on PCBs, which comprises up to 6 SIL radar channels with

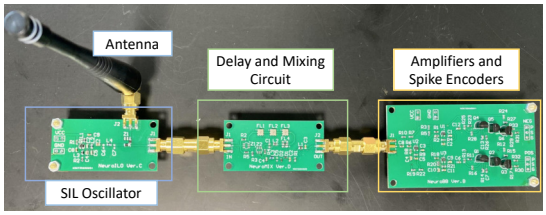

Figure 9: Single-channel SIL radar.

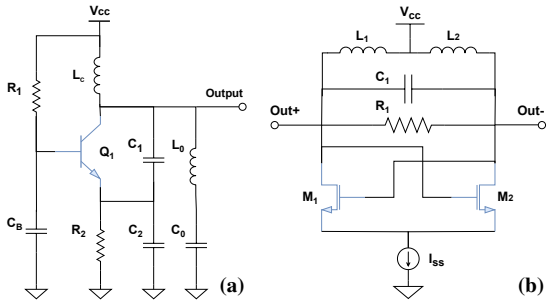

Figure 10: (a) Clapp oscillator, (b) cross-coupled oscillator.

different operating frequencies. We also design a motherboard to enable a robust power distribution to each hardware module.

Ideally, for a real neuromorphic system, the spike encoders should directly interface with neuromorphic computing hardware and send spikes to a pre-trained SNN as input. However, due to a lack of highly specialized neuromorphic processors, we implement the SNN using simulation frameworks that are well-established in the neuromorphic computing research community. Specifically, we adopt the Nengo-DL framework [59] because it supports deep SNN training and accurate emulation of real neuromorphic computing hardware such as Loihi [18]. In addition, due to the need for offline SNN training on the simulated neuromorphic computer, we still need to sample the spikes digitally using an FPGA and store the timestamps on a host PC.

Self-injection locked oscillator. For our prototype, our SILO design employs a Clapp oscillator [62] due to its broad oscillation range and tunability in our targeted UHF band. The oscillator is built using an RF transistor along with discrete LCR components, as shown in Fig. 10a. Infineon BFP620 [81] is selected as the RF transistor for its high transition frequency, which ensures sufficient gain for oscillation at higher operational frequencies. RF inductors and capacitors with high self-resonance frequency are selected to ensure the proper function of the oscillator. The output power of this oscillator design is approximately -20 dBm, and a monopole antenna is used for RF signal transmission. For the IC simulation, we design a cross-coupled oscillator due to its high energy efficiency [62]. It consists of a pair of transistors and features a differential output, as shown in Fig. 10b. The inductor's quality factor substantially impacts energy efficiency, and in the simulation, we use off-chip RF inductors, with their characteristics detailed by the manufacturer [82]. The simulation is carried out using Cadence Virtuoso with a 90 nm generic process design kit.

When an array of SIL sensors is required (Sec. 4.3), their frequencies need to be sufficiently separated to avoid mutual injection

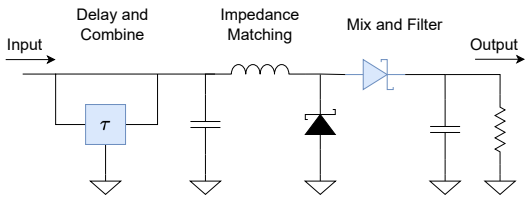

Figure 11: Demodulation circuit schematic.

locking or pulling [61]. The minimum required separation is determined by the quality factor of the LC oscillator and the coupled signal strength. We find empirically that with $\lambda/4$ spacing (Sec. 4.3), the impact of mutual coupling becomes negligible when the center frequencies are separated by at least 10 MHz.

Time delay unit. NEURORADAR requires a time delay to demodulate the motion signal (Sec. 4.2). A longer time delay with minimal insertion loss is desired to attain a large demodulation gain, so the goal is to maximize the gain delay product. We choose to use Surface Acoustic Wave (SAW) filters (RFMi SF2098E [63]) to implement the time delay because of their compact size, low attenuation (1.3dB), and reasonable delay time (28ns). SAW filters operate by converting electrical signals into mechanical vibrations and back into electrical signals. As mechanical or acoustic waves propagate significantly slower than electromagnetic waves, a time delay is introduced when a signal passes through a SAW filter. Notably, the bandwidth of the selected SAW filter is 20–30 MHz, while the maximum observed frequency shift of the oscillator is only 100s of kHz. Therefore, the design ensures that the frequency does not shift outside the bandwidth, where the SAW filters exhibit large attenuation. Moreover, just like regular bandpass filters [52], multiple SAW filters can be cascaded to increase the gain delay product. Given the delay (τ_{SAW}) and gain (G_{SAW}) of a single-stage SAW filter, the optimal number of cascaded SAW filters n can be determined by maximizing the product $n\tau_{SAW}G_{SAW}^{n}$. In the actual implementation, we find that cascading 3 to 4 SAW filters yields the strongest baseband signal.

Low-drive mixer. The mixer multiplies the oscillator signal with its delayed replica and extracts the frequency shift during the time delay. For power efficiency, we opt for passive diode mixers which, in principle, exploit the nonlinearity of diodes to achieve signal mixing. While complex single-balanced and double-balanced mixers offer improved linearity and isolation, they need a high local oscillator (LO) power to drive the diodes [49]. We instead choose a simple single-diode mixer implemented using a low-barrier Schottky diode (Infineon BAT63 [80]), which operates efficiently at ultra-low drive power. To maintain impedance consistency, we shunt another identical diode to the ground to absorb the negative cycle of the input signal, which can only pass the positive cycle through the mixer. Further, a C-L impedance matching circuit precedes the mixer to maximize power delivery from the SILO, and a single-stage LC low-pass filter follows the mixer to attenuate high-frequency signals, as shown in Fig. 11. The demodulated baseband signal has a much lower amplitude than the dynamic range of the spike encoding circuit. Therefore, we need to add an operational amplifier (based on TI TLV521 [34]) module to boost the signal level.

Spike encoding and sampling. We design spike encoder circuits to emulate a LIF neuron, following the description in Sec. 5.1, and

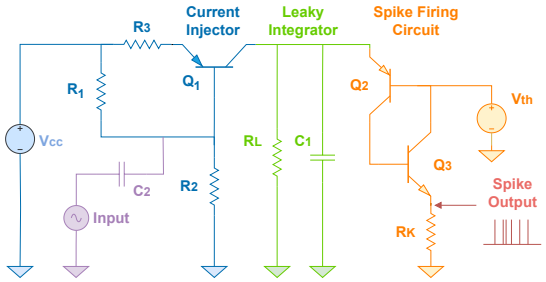

Figure 12: Spike encoding circuit.

Fig. 12. Q_1 (2N3906 [23]) serves as the current injector of the "neuron" that charges the capacitor C_1 with current $I_c(t) \approx V_{in}(t)/R_3$. Q_2 (2N3906) and Q_3 (2N3904 [22]) implement the spike firing circuit, with the firing threshold V_f set by V_{th}. When the voltage of the capacitor V_{C_1} exceeds the firing threshold V_f, Q_2 and Q_3 are turned on shortly to fire a spike and discharge C_1 through R_k. The leaky resistor R_L is directly connected to the ground, setting u_r to zero. To ensure event-driven output, the input signal is AC-coupled, and two instances of LIF encoders are used to represent the negative and positive parts of the signal, respectively. This ensures no spike will be generated if the input signal is constant.

In order to provide spike samples for the Nengo-DL emulation, we connect all output channels of the spike encoder to a lightweight FPGA—Xilinx CMOD-A7 [21]—for sampling. The FPGA samples the spike sequences by polling the digital I/Os. Whenever a spike is detected, each sampling channel creates a frame that contains the timestamp of the spike and the channel index. To support multi-channel parallel sampling, frames from multiple channels are sequenced together using a first-in-first-out (FIFO) buffering block. The output of the FIFO is connected to a universal asynchronous receiver/transmitter (UART) block, which sends the spike frames to the host PC, where the SNN emulation runs.

Power distribution motherboard for the NeuroRadar array. Given that the aforementioned amplifiers offer high gain to the baseband signal, they make the spiking encoding circuit susceptible to noise because even minor disturbances are amplified and may trigger false spikes. To mitigate this, a motherboard with a dedicated power distribution network is designed to provide each SIL radar channel with a stable power supply and suppress noise along the power supply paths. This is accomplished through the incorporation of Linear low-dropout (LDO) power regulators (TPS7A02 [35]) into each amplifier module and spike encoder, which exhibit excellent noise suppression capability. Additionally, the motherboard's large power plane provides a low resistance path for the supply current, further reducing power supply noise.

7 EVALUATION

Microbenchmark of the SIL oscillator. We carry out experiments in a multipath-rich lab setting to assess the motion modulation capability of the self-injection oscillator. To control the experimental conditions, we use a 20cm × 20cm aluminum sheet as a representative target, place it at predefined locations, and measure the SILO's frequency using a spectrum analyzer. As Fig.13a illustrates, the measured frequency shift aligns closely with the theoretical pattern described by Eq.(1). These results indicate that *the frequency of the SILO can indeed be effectively modulated by the movement of nearby reflectors*, and the frequency shift pattern is *not affected by the presence of multipath clutter*.

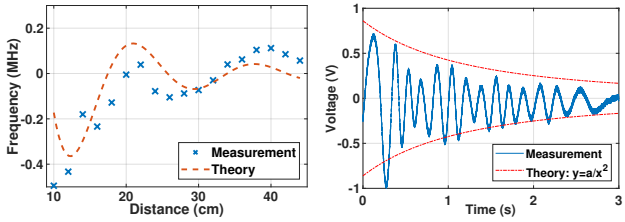

(a) Frequency shift (b) Baseband signal

Figure 13: Modulation and demodulation of a SILO.

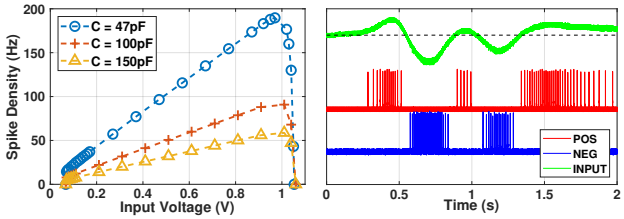

(a) Spike density (b) Spike representation

Figure 14: Characters of the spike encoding circuit.

Microbenchmark of the motion demodulation circuit. We proceed to validate the motion demodulation circuit by attaching it to a running SILO. A target (an adult) moves away from the radar at an approximately constant speed from 0.5m to 3m. Fig. 13b shows the demodulated signal, which follows the sinusoid pattern consistent with Eq. (2). As the target moves away, the reflected signal becomes weaker, causing less frequency variation in the oscillator. The overall baseband signal strength decays approximately proportionally to $1/r^2$. In addition, it can be estimated that the distance the target covers $r = n \cdot \lambda/2 = 2.60m$ (close to the ground truth of 2.5 m), where $n = 14$ is the completed cycles, and the wavelength $\lambda = 37.2\ cm$. Therefore, the result shows that *the motion demodulation circuit can effectively convert the frequency shift of the oscillator into a continuous baseband signal*.

Spike encoder properties. We profile the spike encoding circuit by applying different DC voltage levels at the input and changing the membrane capacitance (Sec. 5.1). Fig. 14a shows the spike density with respect to the input voltage. As analyzed, the spike firing rate can be increased with smaller membrane capacitance. From $V_{in} = 0$ to $0.07V$, the spike encoder is in the dead zone and produces no spikes. When $0.07V < v_{in} < 0.90V$, the spike firing rate increases approximately linearly with the input voltage, as delineated in Eq. (8). When $v_{in} > 0.90V$, the spike rate starts dropping quickly down to 0. The spike density plot shows that *the spike encoding circuit achieves a one-on-one mapping between the input voltage and spike density*. We then connect a real baseband signal from the motion demodulation circuit into the spike encoders and convert it to spike trains. Fig. 14b shows the spike representation of the signal. We find that *the spike generation is indeed event-driven and asynchronous*, as no spike is generated when the input is 0 V, and the spikes can be fired at any random time without any explicit synchronization signal.

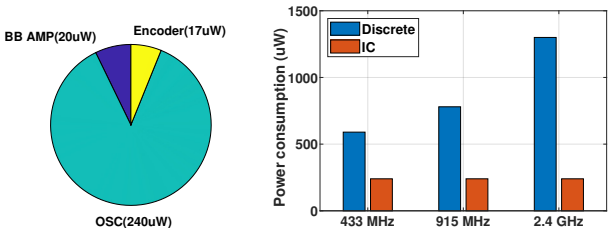

(a) Power breakdown (b) Oscillator power consumption

Figure 15: System power consumption.

Power consumption of the front-end. Next, we characterize the power consumption of a single-channel SIL radar. The radar front-end comprises three main parts: the oscillator, the baseband amplifier, and the spike encoder. As shown in Fig. 15a, the system's power consumption is dominated by the oscillator, which is the sole active RF component in the system. Due to the low signal bandwidth, the baseband amplifier can be designed with low-power consumption, consuming merely 20 µW power. The power consumption of the spike encoder is primarily due to the quiescent current induced by resistor dividers that are used to provide a DC bias. Each spike generation only consumes around 90 pJ and therefore has a negligible impact on the total power consumption. Fig. 15b shows the power consumption of the oscillator at popular operating frequencies in the UHF band. The IC version, which adopts a more power-efficient oscillator structure (discussed in Sec. 6), consumes less power than the discrete version. The total power consumption of the radar front-end falls below 300 µW, underscoring NEURORADAR*'s low-power operational capacity across different operating frequencies in the UHF band.*

8 CASE STUDIES

In this section, we implement and evaluate two use cases based on NEURORADAR: hand gesture recognition and moving target localization. For each case, we train and test the SNN, collect spike chain data, and use them to drive the neuromorphic processor emulation in the NengoDL [59].

A model-based method is employed to estimate the power consumption of the signal processing units. The energy consumption of running an SNN can be calculated as follows:

$$E_{SNN} = N_n * E_n * N_{steps} + E_s * N_{spikes} \quad (10)$$

Here, E_n is the energy consumption of updating the status of a neuron, which must be done for all neurons in the SNN for each emulation timestep (1 ms). N_n is the number of neurons, and N_{steps} is the number of timesteps. E_s denotes the energy consumption of a synaptic operation, which includes generating a spike and passing it to other neurons through synapses. We assume that all the SNNs in our comparisons run on the Intel Loihi neuromorphic chip [18], where $E_s = 23.6pJ$ and $E_n = 81pJ$.

The energy consumption of a conventional ANN is calculated by summing over all the multiply-and-accumulate (MAC) units:

$$E_{ANN} = E_{MAC} * N_{MAC} \quad (11)$$

We assume that all the ANNs run on the Nvidia GTX Titan Black GPU, which has a per-MAC energy consumption of E_{MAC}=3.584 GMAC/W [19].

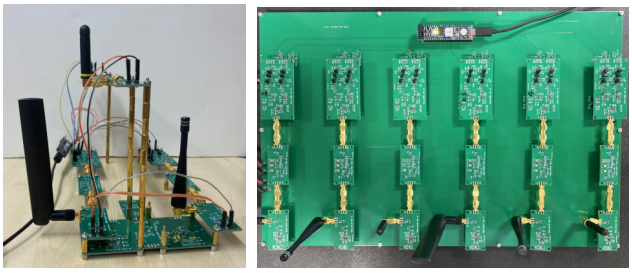

(a) Three-channel setup (b) Six-channel setup

Figure 16: Different NEURORADAR configurations for (a) gesture recognition, and (b) localization.

Radar signal preprocessing is performed using DSP units. Since FFTs are widely used, we employ energy consumption data from the Texas Instruments (TI) TMS320VC5505 [50] to estimate the energy consumption. For other preprocessing algorithms such as digital filtering, we use the MAC energy consumption data of the TI TMS320C6678 (0.853 GMAC/W) [19]. These two TI chips are widely used and represent the state-of-art of DSP unit performance.

8.1 Gesture Recognition

We customize NEURORADAR for hand gesture recognition with two setups, one with a single SIL channel at 915 MHz, and the other with three SIL channels with distinct frequencies around the 866/915 MHz band. For the three-channel setup, two antennas are placed on a horizontal line, with a spacing of $\lambda/4$, with λ being the average carrier wavelength; and the third antenna is placed above the horizontal line and forms an equilateral triangle with the other two antennas, as shown in Fig. 16a. The gestures are made in front of the antenna plane, facing the center of the triangle. The spacing is designed to be comparable to the displacement of a hand when making gestures, affording more distinguishable signal patterns. The elevated antenna provides richer information for vertical hand movement (such as "swipe up" or "swipe down").

Similar to state-of-the-art gesture recognition radar such as Google Soli [41], we define a set of 12 gestures, as shown in Fig. 17. In the gesture set, hand movement direction is diverse within a 3D space (*e.g.* push, pull, left, right, up, and down), and some gestures require 2 hands to move simultaneously. The three-channel version is employed to recognize all 12 gestures. However, the one-channel setup only employs one antenna, which is unable to acquire any angular information. We, therefore, select 4 out of 12 gestures (5, 6, 7, 12) that primarily induce a distinguishable pattern in the range domain, for the one-channel NEURORADAR to recognize.

As each SIL radar channel is paired with two spike encoders (Sec. 6), the three-channel setup produces six spike sequences, while the one-channel setup produces two, and the timestamps of the spikes are recorded for SNN training. Each gesture sample contains sequences of spikes across a time length of 1.5s. As explained in Sec. 5.2, we buffer the input spike sequences, concatenate the output of the delay buffers, and present them together to the convolution layers. We set $n_{dly} = 6$ (6 ms), which results in an input dimension of 6×250 for the three-channel setup and 2×250 for the one-channel setup. In total, we collected 2400 samples, 200 for each gesture. We divide the samples into a training set (1920 samples) and a test set

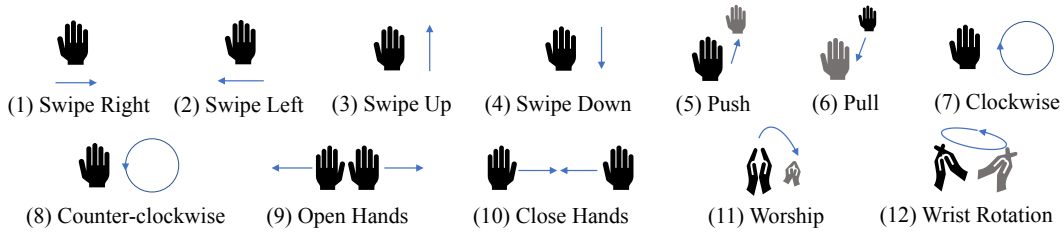

Figure 17: Gesture set definition. Gestures (1) - (8) are single-hand; (9) - (12) are double-hand.

Table 1: SNN specification for gesture recognition.

	type	dimension	channels	kernel	stride	pool
1	conv.	6×250	1	1×16	1×4	1×1
2	conv.	6×59	32	3×16	1×2	1×1
3	conv.	4×22	32	4×8	1×1	1×1
4	dense	1×15	48	N/A	N/A	N/A
5	dense	180	N/A	N/A	N/A	N/A
6	dense	12	N/A	N/A	N/A	N/A

(480 samples) with a random 80/20 split. Although each sample has a fixed length, the starting time of the gesture action is random. We thus perform data augmentation by time-shifting the samples by a small random amount (~50ms). This effectively doubles the training set to 3840 samples.

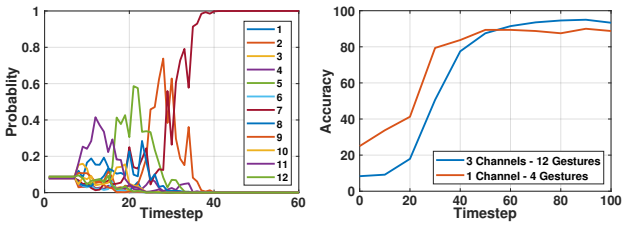

(a) Filtered spike output **(b) Recognition accuracy**

Figure 18: Gesture recognition results.

Fig. 18 summarizes the gesture recognition outcome. As shown in Fig. 18a, the filtered spike signal at each output neuron is interpreted as the probability of each class, and it needs sufficient time steps to stabilize. Fig. 18b indicates that the SNN needs approximately 80 timesteps to produce recognition results with an accuracy exceeding 90%. This means that after gesture operation input is complete, the SNN needs a mere 80 ms to produce a reliable result, which is sufficient for most applications. The confusion matrix in Table 2 shows that the 3-channel NeuroRadar is able to distinguish 12 different gestures, with gesture #10 slightly less accurate than others. Due to the sparsity of spikes, the neuron energy, which linearly increases with emulation time (shown in Eq. (10)), is dominating the overall energy consumption per inference. The number of time steps strikes a tradeoff between higher accuracy and lower energy consumption.

We present a comparison of NeuroRadar with other RF-based gesture recognition systems in Table 3, including works that utilize conventional ANN [68, 77], SNN [6, 7, 67], and other simple machine-learning models [85]. When calculating the power consumption of the RF front-end, duty cycling is considered according to the descriptions provided in the corresponding publications.

Compared with multi-RF-chain radar systems [6, 7, 68, 77], NeuroRadar demonstrates *comparable gesture recognition capabilities*. In comparison to ANN-based systems [68, 77], NeuroRadar

Table 2: Confusion matrix for gesture recognition.

	1	2	3	4	5	6	7	8	9	10	11	12
1	0.94	0	0	0	0	0	0	0	0.03	0	0	0.03
2	0	0.95	0	0	0	0	0	0	0	0	0	0.05
3	0	0	0.93	0	0	0	0	0.02	0.05	0	0	0
4	0	0.03	0	0.91	0	0	0	0.03	0	0	0	0.03
5	0	0	0	0	0.97	0	0	0	0.03	0	0	0
6	0	0.04	0	0	0	0.96	0	0	0	0	0	0
7	0	0	0	0	0	0	1.00	0	0	0	0	0
8	0.02	0	0	0	0	0	0	0.98	0	0	0	0
9	0.02	0	0.04	0	0	0	0	0	0.92	0	0	0.02
10	0.02	0.07	0	0.02	0	0.03	0	0.03	0	0.83	0	0
11	0	0	0	0	0	0	0	0	0	0	0.98	0.02
12	0	0	0	0	0	0	0.02	0	0	0	0	0.98

achieves 3 orders of magnitude reduction in end-to-end power consumption. Other SNN-based gesture recognition systems [6, 7] still rely on conventional computing units (*i.e.*, CPU or DSP) for radar signal pre-processing. This factor dominates the signal processing power and diminishes the benefits of using an SNN. Owing to its full-SNN architecture, NeuroRadar *achieves a power consumption reduction between 78% – 93%* in terms of signal processing (pre-processing and SNN). In addition, compared with other single-channel radar gesture recognition systems with similar gesture sets [67, 85], the single-channel NeuroRadar still reduces the power consumption by at least one order of magnitude.

8.2 Moving Target Localization

To localize a single moving target with an acceptable level of ambiguity, NeuroRadar employs a 6-sensor array with a $\lambda/4$ spacing and diverse carrier frequencies, as simulated in Sec. 4.3. Since NeuroRadar can only detect moving targets, we ask a volunteer to

Figure 19: The area of interest for target location. The maximum distance from the radar is around 6 meters.

Table 3: Comparison of radar systems for gesture recognition.

Work	Front-end	Band (GHz)	RF chains	Pre-proc.	CPU/ GPU	Network/ Algorithm.[1]	Ges-tures	Acc. (%)	P_Front-end (mW)	P_Pre-proc. (µW)	P_ANN (µW)	P_Total (mW)
Sun *et al.* [77]	FMCW	58-63	1Tx+3Rx	FFTs	Yes	ANN-conv.	12	95.79	150.1	498.4	685150	835.74
Tinyradarnn [68]	FMCW	~60	2Tx+2Rx	FFTs	Yes	ANN-conv.	11	92.39	190	642.6	28948	219.59
RadarSNN [7]	FMCW	58-63	1Tx+3Rx	FFTs	Yes	SNN-conv.	8	99.50	9.88	622.2	1	10.51
Arsalan *et al.* [6]	FMCW	58-63	1Tx+3Rx	Filtering	Yes	SNN	8	98.70	5.2	257.4	39	5.49
Safa *et al.* [67]	FMCW	8	1Tx+1Rx	FFTs	Yes	SNN-conv.	5	93.00	0.68	1349.1	20	2.05
Wan *et al.* [85]	Doppler	2.4	1Tx+1Rx	FFTs	Yes	kNN	3	96.70	412.5	154.6	~ 0	412.65
NeuroRadar	SIL	0.915	1Tx/Rx	N/A	No	SNN-conv.	4	90.00	0.28	0	12	0.292
NeuroRadar	SIL	0.915	3Tx/Rx	N/A	No	SNN-conv.	12	94.58	0.83	0	65	0.840

[1] conv. means the neural network comprises convolution layers.

Table 4: SNN specification for localization.

	type	dimension	channels	kernel	stride	pool
1	conv.	12 × 500	1	2 × 32	2 × 1	1 × 4
2	conv.	6 × 118	64	3 × 24	1 × 1	1 × 2
3	conv.	4 × 49	64	4 × 16	1 × 1	1 × 2
4	dense	1 × 17	96	N/A	N/A	N/A
5	dense	1 × 300	N/A	N/A	N/A	N/A
6	dense	1 × 16	N/A	N/A	N/A	N/A

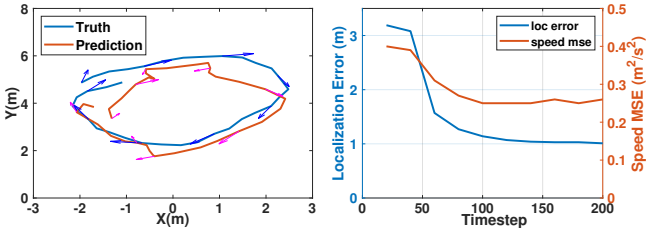

(a) Predicted trace and speed (b) Localization error

Figure 20: NEURORADAR localization result.

walk randomly within the radar's field of view. The maximum distance of the target is around 6m and the angle of view is about 90°, as shown in Fig. 19. We employ a ZED-2i [74] depth camera to obtain the ground-truth location and speed. We collect 6 segments of 10-minute ($600s \times 6 = 3600s$) continuous data for training and testing. For each segment, we allocate the first 480*s* (80%) of data as training samples, reserving the last 120*s* (20%) as test samples. This approach helps to mitigate any inconsistencies that might arise between segments, such as issues with frame alignment with the ground truth. We then further segment the continuous data into 2*s* short frames with a 75% overlap, and each of the short frames becomes a training/test sample. This results in a total of 5742 training samples and 1422 test samples.

Again, the input spike trains are buffered (with $n_{dly} = 4$) and presented collectively to the SNN, resulting in an input dimension of 12 × 500. From each frame, we evenly selected four data points, yielding four sets of location and velocity data: (x_1, y_1, u_1, v_1), (x_2, y_2, u_2, v_2), $\cdots$, (x_4, y_4, u_4, v_4). These sets were used as labels for the regression problem, thus making the output dimension of the neural network 1×16. The specific structure of the SNN model is outlined in Table 4.

Fig. 20a shows the localization result by combining the output of consecutive frames. Similar to the gesture recognition use case, the SNN needs to run for enough timesteps to yield a reasonable result.

Table 5: Indoor tracking system performance comparison.

	NEURORADAR	**Doorpler**[36]
Architecture	SIL	Doppler
RF Chains	6Tx/Rx	2Tx+5Rx
Frequency	~915 MHz	2.4/5.8 GHz
Tx Power	-20 dBm	-10 dBm
Coverage	6 m	<5 m
Front-end Power	1.44 mW	59.9 mW
Signal Proc. Power	0.59 mW	6.4 mW

Fig. 20b shows with about 150 timesteps, a localization accuracy of 1m can be achieved. The mean squared error for speed estimation stabilizes at $0.25 m^2/s^2$. The result implies a tracking delay of 150 ms, which is sufficient for our low-velocity indoor applications. Since we are filtering spike sequences to achieve a continuous value, errors are inevitable and the accuracy is impacted.

To showcase the advantages of NEURORADAR, we compare it with a multi-tone (2.4 and 5.8 GHz) Doppler radar system, Doorpler [36], which utilizes a conventional RF front-end architecture (Sec. 4.1) and signal processing method. Doorpler is a radar-based occupancy sensing system that can detect zone crossing events and estimate the direction of movement at zone transition spots (*e.g.* doorways). Similar to NEURORADAR, Doorpler employs an antenna array to acquire angular information of the target and leverages the Doppler effect to infer the moving direction.

Table 5 compares the performance of the two systems. Due to the extra demodulation gain of SIL radar (Sec. 4.2), NEURORADAR provides a more extensive coverage area than Doorpler, even with 10 dB lower Tx power. Due to its simple SIL structure and power-efficient design, NEURORADAR achieves 1-2 orders of magnitude of reduction in front-end power. The combination of FDA design and neural network allows NEURORADAR to obtain more abundant and accurate sensing information. Unlike Doorpler, which merely detects crossing events and their direction, NEURORADAR offers both location and speed estimation. At the same time, SNN processing significantly reduces the computational power, and the end-to-end system power consumption is reduced by 97%.

9 RELATED WORK

Neuromorphic sensors. Neuromorphic sensors mimic the structure and function of mammalian sensing systems to achieve energy-efficient, and real-time sensory data processing [94]. Neuromorphic sensors have been explored to perceive various types of signals,

such as vision [12, 40], tactile [9, 84], auditory [30, 46, 92], chemicals [27, 73], PH [42, 43], *etc.* Neuromorphic sensing devices have shown great potential for personal healthcare monitoring [13, 71], neuroprosthetics [64] and soft robotics [57]. Built on this line of research, NEURORADAR represents a pioneering step in extending neuromorphic sensing into the RF domain.

SNN-based radar signal processing. Recent research has attempted to use SNNs to process radar signals to reduce power consumption and latency. A majority of the proposed systems first preprocess the raw radar samples using FFT based on conventional computing units (*i.e.*, CPU/DSP), convert the intermediate data, such as range-doppler matrix, into spikes, and then employ SNNs to perform tasks like gesture recognition [5, 7, 8, 33, 38, 66, 67]. Other works analyze the feasibility of using SNNs to perform FFT directly on raw radar samples [6, 45, 83]. In [70], Shaaban *et al.* uses an SNN to directly perform classification on time-domain radar frames. However, all these methods are based on traditional radar front-ends, and the radar signal must first be sampled using ADCs, then processed and converted into spikes digitally. In contrast, NEURORADAR adopts a novel front-end that produces spike sequences and can directly interface with energy-efficient neuromorphic computing hardware.

Self-injection locked radar. Wang *et al.* [90] were among the first to analyze the advantages of SIL radar and demonstrate its ability to perform vital sign sensing. Numerous variations of SIL radar have been explored, such as a single-antenna SIL radar [89], mutual injection locked radar [88], and bistatic SIL radar [87]. Tang *et al.* [79] compared the performance of SIL radar with Doppler radar and proposed innovative designs to eliminate null detection points. Hsu *et al.* [32] studies a mutually injection-locked oscillator array and demonstrates that by adjusting the tuning voltages of the oscillators, a beamforming SIL radar can be achieved. In contrast, NEURORADAR studies the sensing capabilities of a frequency-diverse SIL radar array, and performs gesture recognition and target localization.

Low-power radar sensing. The RF amplifiers in radar systems constitute a substantial source of power consumption. Prior research has investigated energy-efficient, high power density amplifiers, such as gallium nitride (GaN) amplifiers [17, 29, 91], to reduce the power consumption of the radar front-end. Radar waveform design represents another approach to power reduction. For instance, impulse radio ultra-wideband (IR-UWB) radars [16, 54] emit short-duration pulses that helps achieve lower power compared to the widely used FMCW radar. Another approach for low-power radar design is utilizing existing radio sources. Passive radars, which utilize existing radio sources such as television [37] and Wi-Fi [25, 39], eliminate the need for transmitting their own signals. Consequently, they provide a cost-effective, low-power alternative to conventional active radar systems. NEURORADAR offers a distinctive solution to this challenge, achieving low power consumption by incorporating neuromorphic engineering into radar sensing.

10 DISCUSSIONS AND FUTURE WORKS

Multi-target localization. While our experiments underscore NEURORADAR's proficiency in sensing a singular moving target, the model can seamlessly be extended for multi-target sensing. With N targets in the vicinity, Eq. (4) can be rewritten as follows:

$$\vec{s}(\vec{l},\vec{v}) = \sum_{n=1}^{N}[a_n \sin(4\pi f_k \frac{r_{k,m,n}}{c})]_{k=0,m=0}^{k=K-1,m=M-1},$$

The observed frequency shift becomes the superposition of the frequency shift induced by each moving target. The potential to localize targets is dependent upon the ambiguity area resulting from solving the optimization problem in Eq. (5). As an increase in targets would induce more ambiguity in the observation, a larger number of sensors are required to distinctly localize each target. Subsequently, following the same single-target paradigm, the SNN can be trained to identify multiple targets based on the observations.

Reducing the form factor. For broader real-world applicability, minimizing the system's form factor is pivotal. Though the majority of the NEURORADAR components can be integrated into an IC, the system's form factor is predominantly dictated by the antenna dimensions and inter-antenna spacing. The monopole antennas we employ, while robust, are size-intensive. Replacements like compact PCB loop antennas, often adopted in low-frequency IoT devices [51], could be more feasible. In addition, the current implementation has restricted the system to the UHF band because of the frequency limitation of the SAW filters (below 3 GHz). Developing NEURORADAR sensors at higher frequencies can substantially reduce its form factor, due to smaller antennas and a more compact inter-antenna spacing. We leave the exploration of such solutions for future work.

Limitations and suitable applications. While the SIL radar's streamlined architecture facilitates energy-efficient operation, its free-running oscillator remains vulnerable to external disturbances. In-band signals from external sources may cause frequency pulling or injection locking [60] to the oscillator, thereby jeopardizing the SIL radar's sensing capability either partially or wholly. Moreover, SIL radars need to be stationary to function because otherwise, all the reflectors in the environment would be moving relatively to the radar, each adding a frequency shift to the oscillator, and making it impractical to perform sensing. Nevertheless, in a controlled indoor setting devoid of interference, NEURORADAR can support various applications such as surveillance systems, vital sign monitors, motion tracking, and gesture-based controls.

11 CONCLUSION

In this work, we have introduced NEURORADAR, a novel and pioneering approach in radar systems that fully embraces the principles of neuromorphic sensing. Through the joint design of analog hardware and spike signal processing, NEURORADAR achieves superior energy efficiency. Through gesture recognition and localization tasks, NEURORADAR has demonstrated its capability while maintaining a power consumption significantly lower than that of traditional radar systems. This research marks a significant step forward, providing a unique and innovative solution for radar sensing in energy-constrained IoT devices.

ACKNOWLEDGMENTS

We appreciate the insightful comments and feedback from the anonymous reviewers and shepherd. The work reported in this paper is supported in part by the NSF under Grants CNS-2312715, CNS-1901048, CNS-1925767, and CNS-2128588.

REFERENCES

[1] Robert Adler. 1946. A study of locking phenomena in oscillators. *Proceedings of the IRE* 34, 6 (1946), 351–357.

[2] Zeeshan Ahmad, Meng Chen, and Shu-Di Bao. 2021. Beam pattern analysis of frequency diverse array radar: a review. *EURASIP Journal on Wireless Communications and Networking* 2021, 1 (2021), 1–33.

[3] Filipp Akopyan, Jun Sawada, Andrew Cassidy, Rodrigo Alvarez-Icaza, John Arthur, Paul Merolla, Nabil Imam, Yutaka Nakamura, Pallab Datta, Gi-Joon Nam, et al. 2015. Truenorth: Design and tool flow of a 65 mw 1 million neuron programmable neurosynaptic chip. *IEEE transactions on computer-aided design of integrated circuits and systems* 34, 10 (2015), 1537–1557.

[4] Arnon Amir, Brian Taba, David Berg, Timothy Melano, Jeffrey McKinstry, Carmelo Di Nolfo, Tapan Nayak, Alexander Andreopoulos, Guillaume Garreau, Marcela Mendoza, et al. 2017. A low power, fully event-based gesture recognition system. In *Proceedings of the IEEE conference on computer vision and pattern recognition*. 7388–7397. https://doi.org/10.1109/CVPR.2017.781

[5] Muhammad Arsalan, Mateusz Chmurski, Avik Santra, Moamen El-Masry, Robert Weigel, and Vadim Issakov. 2021. Resource efficient gesture sensing based on FMCW radar using spiking neural networks. In *2021 IEEE MTT-S International Microwave Symposium (IMS)*. IEEE, 78–81.

[6] Muhammad Arsalan, Avik Santra, and Vadim Issakov. 2022. Power-efficient gesture sensing for edge devices: mimicking fourier transforms with spiking neural networks. *Applied Intelligence* (2022), 1–16.

[7] Muhammad Arsalan, Avik Santra, and Vadim Issakov. 2022. RadarSNN: A resource efficient gesture sensing system based on mm-wave radar. *IEEE Transactions on Microwave Theory and Techniques* 70, 4 (2022), 2451–2461.

[8] Dighanchal Banerjee, Smriti Rani, Arun M George, Arijit Chowdhury, Sounak Dey, Arijit Mukherjee, Tapas Chakravarty, and Arpan Pal. 2020. Application of spiking neural networks for action recognition from radar data. In *2020 International Joint Conference on Neural Networks (IJCNN)*. IEEE, 1–10.

[9] Tom Birkoben, Henning Winterfeld, Simon Fichtner, Adrian Petraru, and Hermann Kohlstedt. 2020. A spiking and adapting tactile sensor for neuromorphic applications. *Scientific reports* 10, 1 (2020), 1–11.

[10] Peter Blouw, Xuan Choo, Eric Hunsberger, and Chris Eliasmith. 2019. Benchmarking keyword spotting efficiency on neuromorphic hardware. In *Proceedings of the 7th annual neuro-inspired computational elements workshop*. 1–8.

[11] Peter Blouw and Chris Eliasmith. 2020. Event-driven signal processing with neuromorphic computing systems. In *ICASSP 2020-2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 8534–8538.

[12] Christian Brandli, Raphael Berner, Minhao Yang, Shih-Chii Liu, and Tobi Delbruck. 2014. A 240× 180 130 db 3 μs latency global shutter spatiotemporal vision sensor. *IEEE Journal of Solid-State Circuits* 49, 10 (2014), 2333–2341.

[13] Karla Burelo, Mohammadali Sharifshazileh, Niklaus Krayenbühl, Georgia Ramantani, Giacomo Indiveri, and Johannes Sarnthein. 2021. A spiking neural network (SNN) for detecting high frequency oscillations (HFOs) in the intraoperative ECoG. *Scientific Reports* 11, 1 (2021), 6719.

[14] Heng-Chia Chang. 2003. Stability analysis of self-injection-locked oscillators. *IEEE Transactions on Microwave Theory and Techniques* 51, 9 (2003).

[15] Du Chen, Yuan Li, Dongming Xu, John G Harris, and José Carlos Principe. 2006. Asynchronous biphasic pulse signal coding and its CMOS realization. In *2006 IEEE International Symposium on Circuits and Systems (ISCAS)*. IEEE, 4–pp.

[16] Zhe Chen, Tianyue Zheng, and Jun Luo. 2021. Octopus: a practical and versatile wideband MIMO sensing platform. In *Proceedings of the 27th Annual International Conference on Mobile Computing and Networking*. 601–614.

[17] Ken B Cooper, Raquel Rodriguez Monje, Robert J Dengler, Corey J Cochrane, Maria Alonso-Delpino, Adrian Tang, Tristan Ossama El Bouayadi, and Omkar Pradhan. 2020. A compact, low power consumption, and highly sensitive 95 GHz Doppler radar. *IEEE Sensors Journal* 20, 11 (2020), 5865–5875.

[18] Mike Davies, Narayan Srinivasa, Tsung-Han Lin, Gautham Chinya, Yongqiang Cao, Sri Harsha Choday, Georgios Dimou, Prasad Joshi, Nabil Imam, Shweta Jain, et al. 2018. Loihi: A neuromorphic manycore processor with on-chip learning. *Ieee Micro* 38, 1 (2018), 82–99.

[19] Brian Degnan, Bo Marr, and Jennifer Hasler. 2015. Assessing trends in performance per watt for signal processing applications. *IEEE Transactions on Very Large Scale Integration (VLSI) Systems* 24, 1 (2015).

[20] Debashis Dhar, Paul TM van Zeijl, Dusan Milosevic, Hao Gao, and Arthur HM van Roermund. 2017. Modeling and analysis of the effects of PLL phase noise on FMCW radar performance. In *2017 IEEE International Symposium on Circuits and Systems (ISCAS)*. IEEE, 1–4.

[21] Digilent. 2020. Cmod A7. https://digilent.com/reference/programmable-logic/cmod-a7/start

[22] NTE Electronics. 2010. 2N3904. https://www.nteinc.com/specs/original/2N3904.pdf

[23] NTE Electronics. 2010. 2N3906. https://www.nteinc.com/specs/original/2N3905_06.pdf

[24] Steve B Furber, Francesco Galluppi, Steve Temple, and Luis A Plana. 2014. The spinnaker project. *Proc. IEEE* 102, 5 (2014), 652–665.

[25] Chuhan Gao, Yilong Li, and Xinyu Zhang. 2018. LiveTag: Sensing Human-Object Interaction through Passive Chipless WiFi Tags. In *USENIX Symposium on Networked Systems Design and Implementation (NSDI)*.

[26] Wulfram Gerstner, Werner M Kistler, Richard Naud, and Liam Paninski. 2014. *Neuronal dynamics: From single neurons to networks and models of cognition*. Cambridge University Press.

[27] Martina Giordani, Marcello Berto, Michele Di Lauro, Carlo A Bortolotti, Michele Zoli, and Fabio Biscarini. 2017. Specific dopamine sensing based on short-term plasticity behavior of a whole organic artificial synapse. *ACS sensors* 2, 12 (2017), 1756–1760.

[28] Jian Gong, Xinyu Zhang, Kaixin Lin, Ju Ren, Yaoxue Zhang, and Wenxun Qiu. 2021. RF vital sign sensing under free body movement. *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies* 5, 3 (2021).

[29] K Husna Hamza and D Nirmal. 2020. A review of GaN HEMT broadband power amplifiers. *AEU-International Journal of Electronics and Communications* 116 (2020), 153040.

[30] Yongli He, Sha Nie, Rui Liu, Shanshan Jiang, Yi Shi, and Qing Wan. 2019. Spatiotemporal information processing emulated by multiterminal neuro-transistor networks. *Advanced Materials* 31, 21 (2019), 1900903.

[31] Miriam Henning, Giordano Ramos-Traslosheros, Burak Gür, and Marion Silies. 2022. Populations of local direction–selective cells encode global motion patterns generated by self-motion. *Science Advances* 8, 3 (2022).

[32] Chung-Yi Hsu, Cho-Ying Chuang, Fu-Kang Wang, Tzyy-Sheng Horng, and Lih-Tyng Hwang. 2017. Detection of vital signs for multiple subjects by using self-injection-locked radar and mutually injection-locked beam scanning array. In *2017 IEEE MTT-S International Microwave Symposium (IMS)*. IEEE, 991–994.

[33] Jiaxin Huang, Bernhard Vogginger, Pascal Gerhards, Felix Kreutz, Florian Kelber, Daniel Scholz, Klaus Knobloch, and Christian Georg Mayr. 2022. Real-time radar gesture classification with spiking neural network on spinnaker 2 prototype. In *2022 IEEE 4th International Conference on Artificial Intelligence Circuits and Systems (AICAS)*. IEEE, 362–365.

[34] Texas Instruments. 2016. TLV521. https://www.ti.com/product/TLV521

[35] Texas Instruments. 2022. TPS7A02. https://www.ti.com/product/TPS7A02

[36] Avinash Kalyanaraman, Elahe Soltanaghaei, and Kamin Whitehouse. 2019. Doorpler: A radar-based system for real-time, low power zone occupancy sensing. In *2019 IEEE Real-Time and Embedded Technology and Applications Symposium (RTAS)*. IEEE, 42–53.

[37] Bryce Kellogg, Vamsi Talla, and Shyamnath Gollakota. 2014. Bringing gesture recognition to all devices. In *11th USENIX Symposium on Networked Systems Design and Implementation (NSDI 14)*. 303–316.

[38] Felix Kreutz, Pascal Gerhards, Bernhard Vogginger, Klaus Knobloch, and Christian Georg Mayr. 2021. Applied spiking neural networks for radar-based gesture recognition. In *2021 7th International Conference on Event-Based Control, Communication, and Signal Processing (EBCCSP)*. IEEE, 1–4.

[39] Krzysztof Kulpa, Stanisław Rzewuski, Zbigniew Gajo, and Mateusz Malanowski. 2011. Concept of multistatic passive radar based on wireless packet communication systems. In *Proceedings of 2011 IEEE CIE International Conference on Radar*, Vol. 1. IEEE, 149–152.

[40] Patrick Lichtsteiner, Christoph Posch, and Tobi Delbruck. 2008. A 128× 128 120 dB 15 μs Latency Asynchronous Temporal Contrast Vision Sensor. *IEEE Journal of Solid-State Circuits* 43, 2 (2008), 566–576. https://doi.org/10.1109/JSSC.2007.914337

[41] Jaime Lien, Nicholas Gillian, M Emre Karagozler, Patrick Amihood, Carsten Schwesig, Erik Olson, Hakim Raja, and Ivan Poupyrev. 2016. Soli: Ubiquitous gesture sensing with millimeter wave radar. *ACM Transactions on Graphics (TOG)* 35, 4 (2016), 1–19.

[42] Ning Liu, Yu Liu, Junping Hu, Yuping He, Xiaohang Zhang, and Qing Wan. 2019. PH-dependent plasticity regulation in proton/electron hybrid oxide-based synaptic transistors. *Applied Surface Science* 481 (2019), 1412–1417.

[43] Ning Liu, Li Qiang Zhu, Ping Feng, Chang Jin Wan, Yang Hui Liu, Yi Shi, and Qing Wan. 2015. Flexible sensory platform based on oxide-based neuromorphic transistors. *Scientific reports* 5, 1 (2015), 1–9.

[44] Shih-Chii Liu, Tobi Delbruck, Giacomo Indiveri, Adrian Whatley, and Rodney Douglas. 2014. *Event-based neuromorphic systems*. John Wiley & Sons.

[45] Javier López-Randulfe, Tobias Duswald, Zhenshan Bing, and Alois Knoll. 2021. Spiking neural network for fourier transform and object detection for automotive radar. *Frontiers in Neurorobotics* 15 (2021), 69.

[46] Richard F Lyon and Carver Mead. 1988. An analog electronic cochlea. *IEEE Transactions on Acoustics, Speech, and Signal Processing* 36, 7 (1988), 1119–1134.

[47] Wolfgang Maass. 1997. Networks of spiking neurons: the third generation of neural network models. *Neural networks* 10, 9 (1997), 1659–1671.

[48] Mobeen Mahmood and Hasan Mir. 2019. FDA transmit beampattern synthesis using piecewise trigonometric frequency offset. *IET Radar, Sonar & Navigation* 13, 7 (2019), 1149–1153.

[49] Ferenc Marki and Christopher Marki. 2010. Mixer basics primer. *Marki Microwave* (2010), 116.

[50] Mark McKeown. 2010. FFT Implementation on the TMS320VC5505, TMS320C5505, and TMS320C5515 DSPs. *Texas Instruments Incorporated, White Paper SPRABB6B* (2010).

[51] Patrick P Mercier, Saurav Bandyopadhyay, Andrew C Lysaght, Konstantina M Stankovic, and Anantha P Chandrakasan. 2014. A sub-nW 2.4 GHz transmitter for low data-rate sensing applications. *IEEE journal of solid-state circuits* 49, 7 (2014), 1463–1474.

[52] Mini-Circuits. 2020. *Advantages of Cascading Reflectionless Filters*. https://www.minicircuits.com/appdoc/AN75-008.html

[53] Saber Moradi, Ning Qiao, Fabio Stefanini, and Giacomo Indiveri. 2017. A scalable multicore architecture with heterogeneous memory structures for dynamic neuromorphic asynchronous processors (DYNAPs). *IEEE transactions on biomedical circuits and systems* 12, 1 (2017), 106–122.

[54] Van-Han Nguyen and Jae-Young Pyun. 2015. Location detection and tracking of moving targets by a 2D IR-UWB radar system. *sensors* 15, 3 (2015), 6740–6762.

[55] John Nolan, Kun Qian, and Xinyu Zhang. 2021. RoS: passive smart surface for roadside-to-vehicle communication. In *Proceedings of the ACM SIGCOMM Conference*.

[56] Byung-Kwon Park, Olga Boric-Lubecke, and Victor M Lubecke. 2007. Arctangent demodulation with DC offset compensation in quadrature Doppler radar receiver systems. *IEEE transactions on Microwave theory and techniques* 55, 5 (2007), 1073–1079.

[57] Hea-Lim Park, Yeongjun Lee, Naryung Kim, Dae-Gyo Seo, Gyeong-Tak Go, and Tae-Woo Lee. 2020. Flexible neuromorphic electronics for computing, soft robotics, and neuroprosthetics. *Advanced Materials* 32, 15 (2020), 1903558.

[58] Kun Qian, Zhaoyuan He, and Xinyu Zhang. 2020. 3D point cloud generation with millimeter-wave radar. *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies* (2020).

[59] Daniel Rasmussen. 2019. NengoDL: Combining deep learning and neuromorphic modelling methods. *Neuroinformatics* 17, 4 (2019), 611–628.

[60] Behzad Razavi. 2004. A study of injection locking and pulling in oscillators. *IEEE journal of solid-state circuits* 39, 9 (2004), 1415–1424.

[61] Behzad Razavi. 2006. Mutual injection pulling between oscillators. In *IEEE Custom Integrated Circuits Conference 2006*. IEEE, 675–678.

[62] Behzad Razavi and Razavi Behzad. 2012. *RF microelectronics*. Vol. 2. Prentice hall New York.

[63] RFMi. 2022. SF2098E. https://www.rfmi.co/home/products/SAW_BAW

[64] Udaya Bhaskar Rongala, Alberto Mazzoni, and Calogero Maria Oddo. 2015. Neuromorphic artificial touch for categorization of naturalistic textures. *IEEE transactions on neural networks and learning systems* 28, 4 (2015), 819–829.

[65] Bodo Rueckauer, Iulia-Alexandra Lungu, Yuhuang Hu, Michael Pfeiffer, and Shih-Chii Liu. 2017. Conversion of continuous-valued deep networks to efficient event-driven networks for image classification. *Frontiers in neuroscience* 11 (2017), 682.

[66] Ali Safa, Andre Bourdoux, Ilja Ocket, Francky Catthoor, and Georges GE Gielen. 2021. On the use of spiking neural networks for ultralow-power radar gesture recognition. *IEEE Microwave and Wireless Components Letters* 32, 3 (2021), 222–225.

[67] Ali Safa, Federico Corradi, Lars Keuninckx, Ilja Ocket, André Bourdoux, Francky Catthoor, and Georges GE Gielen. 2021. Improving the accuracy of spiking neural networks for radar gesture recognition through preprocessing. *IEEE Transactions on Neural Networks and Learning Systems* (2021).

[68] Moritz Scherer, Michele Magno, Jonas Erb, Philipp Mayer, Manuel Eggimann, and Luca Benini. 2021. Tinyradarnn: Combining spatial and temporal convolutional neural networks for embedded gesture recognition with short range radars. *IEEE Internet of Things Journal* 8, 13 (2021), 10336–10346.

[69] Abhronil Sengupta, Yuting Ye, Robert Wang, Chiao Liu, and Kaushik Roy. 2019. Going Deeper in Spiking Neural Networks: VGG and Residual Architectures. arXiv:1802.02627

[70] Ahmed Shaaban, Wolfgang Furtner, Robert Weigel, and Fabian Lurz. 2022. Spiking Neural Networks for Gesture Recognition Using Time Domain Radar Data. In *2022 19th European Radar Conference (EuRAD)*. IEEE, 33–36.

[71] Mohammadali Sharifshazileh, Karla Burelo, Johannes Sarnthein, and Giacomo Indiveri. 2021. An electronic neuromorphic system for real-time detection of high frequency oscillations (HFO) in intracranial EEG. *Nature communications* 12, 1 (2021), 1–14.

[72] Kashif Siddiq, Robert J Watson, Steve R Pennock, Philip Avery, Richard Poulton, and Ben Dakin-Norris. 2015. Phase noise analysis in FMCW radar systems. In *2015 European Microwave Conference (EuMC)*. IEEE, 1523–1526.

[73] Zhiqi Song, Yanhong Tong, Xiaoli Zhao, Hang Ren, Qingxin Tang, and Yichun Liu. 2019. A flexible conformable artificial organ-damage memory system towards hazardous gas leakage based on a single organic transistor. *Materials Horizons* 6, 4 (2019), 717–726.

[74] Stereolabs. 2023. ZED-2i. https://www.stereolabs.com/zed-2i/

[75] Terrence C Stewart. 2012. A technical overview of the neural engineering framework. *University of Waterloo* 110 (2012).

[76] Jan Stuijt, Manolis Sifalakis, Amirreza Yousefzadeh, and Federico Corradi. 2021. μBrain: An event-driven and fully synthesizable architecture for spiking neural networks. *Frontiers in neuroscience* (2021), 538.

[77] Yuliang Sun, Tai Fei, Xibo Li, Alexander Warnecke, Ernst Warsitz, and Nils Pohl. 2020. Real-time radar-based gesture detection and recognition built in an edge-computing platform. *IEEE Sensors Journal* 20, 18 (2020), 10706–10716.

[78] Dwi Joko Suroso, Deepak Gautam, and Sunarno Sunarno. 2019. Spatial aliasing effects on beamforming performance in large-spacing antenna array. *Communications in Science and Technology* 4, 1 (2019), 1–6.

[79] Dongyang Tang, Davi VQ Rodrigues, Michael C Brown, and Changzhi Li. 2022. Dual Null Detection Points Removal and Time-Domain Sensitivity Analysis of a Self-Injection-Locked Radar for Small-Amplitude Motion Sensing. *IEEE Transactions on Microwave Theory and Techniques* 70, 9 (2022), 4263–4272.

[80] Infineon Technologies. 2011. BAT63. https://www.infineon.com/cms/en/product/rf/rf-diode/rf-mixer-and-detector-schottky-diode/bat63-02v/

[81] Infineon Technologies. 2019. BFP620. https://www.infineon.com/cms/en/product/rf/rf-transistor/low-noise-rf-transistors/bfp620/

[82] Johanson Technology. 2020. Wirewound inductors. https://www.johansontechnology.com/wirewound-inductors

[83] Bernhard Vogginger, Felix Kreutz, Javier López-Randulfe, Chen Liu, Robin Dietrich, Hector A Gonzalez, Daniel Scholz, Nico Reeb, Daniel Auge, Julian Hille, et al. 2022. Automotive Radar Processing With Spiking Neural Networks: Concepts and Challenges. *Frontiers in neuroscience* 16 (2022).

[84] Changjin Wan, Geng Chen, Yangming Fu, Ming Wang, Naoji Matsuhisa, Shaowu Pan, Liang Pan, Hui Yang, Qing Wan, Liqiang Zhu, et al. 2018. An artificial sensory neuron with tactile perceptual learning. *Advanced Materials* 30, 30 (2018), 1801291.

[85] Qian Wan, Yiran Li, Changzhi Li, and Ranadip Pal. 2014. Gesture recognition for smart home applications using portable radar sensors. In *2014 36th annual international conference of the IEEE engineering in medicine and biology society*. IEEE, 6414–6417.

[86] Ban Wang, Gabriele Tasselli, Cyril Botteron, and Pierre-André Farine. 2014. Building Blocks for a 24 GHz Phased-Array Front-End in CMOS Technology for Smart Streetlights. *Progress In Electromagnetics Research B* 61, ARTICLE (2014), 99–119.

[87] Fu-Kang Wang, You-Rung Chou, Yen-Chen Chiu, and Tzyy-Sheng Horng. 2015. Chest-worn health monitor based on a bistatic self-injection-locked radar. *IEEE Transactions on Biomedical Engineering* 62, 12 (2015), 2931–2940.

[88] Fu-Kang Wang, Tzyy-Sheng Horng, Kang-Chun Peng, Je-Kuan Jau, Jian-Yu Li, and Cheng-Chung Chen. 2011. Single-antenna Doppler radars using self and mutual injection locking for vital sign detection with random body movement cancellation. *IEEE Transactions on Microwave Theory and Techniques* 59, 12 (2011), 3577–3587.

[89] Fu-Kang Wang, Tzyy-Sheng Horng, Kang-Chun Peng, Je-Kuan Jau, Jian-Yu Li, and Cheng-Chung Chen. 2011. Single-antenna Doppler radars using self and mutual injection locking for vital sign detection with random body movement cancellation. *IEEE Transactions on Microwave Theory and Techniques* 59, 12 (2011), 3577–3587.

[90] Fu-Kang Wang, Chien-Jung Li, Chieh-Hsun Hsiao, Tzyy-Sheng Horng, Jenshan Lin, Kang-Chun Peng, Je-Kuan Jau, Jian-Yu Li, and Cheng-Chung Chen. 2010. A Novel Vital-Sign Sensor Based on a Self-Injection-Locked Oscillator. *IEEE Transactions on Microwave Theory and Techniques* 58, 12 (2010), 4112–4120. https://doi.org/10.1109/TMTT.2010.2087349

[91] Colin S Whelan, Nicholas J Kolias, Steven Brierley, Chris MacDonald, and Steven Bernstein. 2012. GaN technology for radars. In *2012 CS MANTECH Conference, Boston, MA April*.

[92] Minhao Yang, Chen-Han Chien, Tobi Delbruck, and Shih-Chii Liu. 2016. A 0.5 V 55μW 64×2 Channel Binaural Silicon Cochlea for Event-Driven Stereo-Audio Sensing. *IEEE Journal of Solid-State Circuits* 51, 11 (2016), 2554–2569.

[93] Yole. 2020. Status of the Radar Industry: Players, Applications and Technology Trends 2020. (2020). https://www.yolegroup.com/product/report/status-of-the-radar-industry-players-applications-and-technology-trends-2020/

[94] Mingyue Zeng, Yongli He, Chenxi Zhang, and Qing Wan. 2021. Neuromorphic devices for bionic sensing and perception. *Frontiers in Neuroscience* (2021), 805.