RippleSense: Scalable and Efficient Wideband Spectrum Sensing

Andreas Kuster
College of Computing and Data Science
Nanyang Technological University
Singapore
S220003@e.ntu.edu.sg

Yanbo Zhang
College of Computing and Data Science
Nanyang Technological University
Singapore
yanbo001@e.ntu.edu.sg

Mo Li
Computer Science and Engineering
Hong Kong University of Science and Technology
Hong Kong
lim@cse.ust.hk

Abstract

Dynamic Spectrum Sharing (DSS) is essential for optimizing spectrum utilization in modern wireless systems, but it requires high analog-to-digital converter (ADC) sampling rates, leading to increased costs and power consumption. Existing sub-Nyquist sampling techniques partially address these issues but struggle with dense spectrum adaptability, real-time performance, and efficiency. This paper presents RippleSense, a scalable and efficient wideband spectrum sensing approach capable of capturing GHz of densely occupied spectrum at sub-Nyquist ADC sampling rates. A novel sub-Nyquist sampling method is introduced by injecting distinct signatures into observed signals over different Nyquist zones before sampling, allowing programmatically reconstructing the full spectrum even after the Nyquist zones are folded to baseband due to inadequate ADC sampling rates. To showcase the scalability of this approach, a high-performance, multi-GHz spectrum sensing platform is implemented together with a highly parallelizable reconstruction algorithm that can process the data stream in real-time. Experimental evaluation has shown that the proposed approach supports operating configurations with signal-to-noise ratios as low as −10 dB and time resolutions down to 10 ns, enabling the capture of single radar pulses across bandwidths of up to 10 GHz using our prototype.

CCS Concepts

• **Hardware** → **Digital signal processing**; **Wireless devices**; **Reconfigurable logic applications**; • **Networks** → **Wireless local area networks**; **Network monitoring**.

Keywords

Wireless communication, spectrum sensing, signal processing, wideband, signal injection

ACM Reference Format:
Andreas Kuster, Yanbo Zhang, and Mo Li. 2026. RippleSense: Scalable and Efficient Wideband Spectrum Sensing. In *ACM/IEEE International Conference on Embedded Artificial Intelligence and Sensing Systems (SenSys '26), May 11–14, 2026, Saint Malo, France.* ACM, New York, NY, USA, 13 pages. https://doi.org/10.1145/3774906.3800487



This work is licensed under a Creative Commons Attribution 4.0 International License.

SenSys '26, Saint Malo, France

© 2026 Copyright held by the owner/author(s).

ACM ISBN 979-8-4007-2309-4/26/05
https://doi.org/10.1145/3774906.3800487

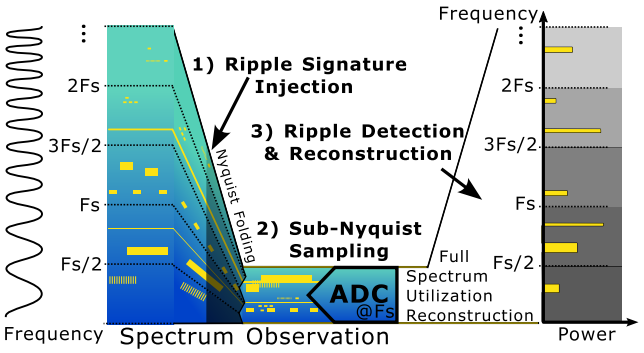


Figure 1: Dynamic spectrum sharing using sub-Nyquist sampling (Fs=sampling rate) by injecting a distinct, per-Nyquist zone signature, in combination with an efficient signature detection and reconstruction signal processing algorithm.

1 Introduction

The scarcity of RF spectrum resources has become a major limitation for advancing new technologies of wireless communication. Dynamic Spectrum Sharing (DSS) has emerged as a key solution, enabling real-time detection and utilization of underutilized spectrum bands, thereby enhancing the flexibility and efficiency of modern communication systems. This technology has been successfully applied in various scenarios, such as supporting IEEE 802.22 and IEEE 802.11af standards in TV White Spaces [6] to provide broadband access to rural areas, facilitating LTE and 5G network expansion through the tiered sharing model of the 3.5 GHz Citizens Broadband Radio Service (CBRS) [7], and reducing interference in the 6 GHz band for Wi-Fi 6E [8] with Automated Frequency Coordination (AFC). Additionally, DSS has demonstrated its versatility in applications like smart grid communications (e.g., IEEE 802.15.4g), and drone networks. However, implementing efficient DSS requires unlicensed devices to perform real-time wideband spectrum sensing across a broad frequency range to identify and utilize temporarily unoccupied channels. This capability demands high-speed ADCs capable of sampling at the Nyquist rate, which poses significant technical and cost barriers.

To better understand this problem, we conducted a large-scale survey of 9134 ADCs[1], analyzing the trade-offs between sampling rate, power dissipation, price, and power efficiency (Figure 2). By applying a linear regression model (95% CI, 1000x bootstrapped), we found that power dissipation and price are strongly correlated with

[1]The data source is the Mouser Electronics online inventory.

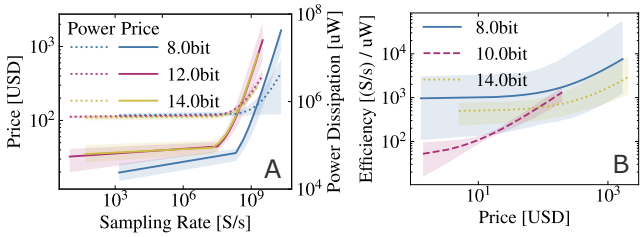


Figure 2: Large-scale survey (9134 ADCs) on sampling rate F_s, power dissipation, price and power efficiency.

the sampling rate for a given resolution class (Figure 2A). Furthermore, more energy-efficient chips tend to be more expensive (Figure 2B). The results reveal that higher sampling rates and increased resolution generally lead to increased power dissipation and cost, highlighting the difficulties in achieving optimal performance for DSS applications.

Researchers have attempted to perform wideband spectrum sensing without adhering to the Nyquist rate [2, 5, 10–13, 15, 17, 20, 27, 28, 31, 36, 42] to mitigate this issue. However, the existing methods face limitations that hinder their practical adoption: (1) *Inadequate adaptation to dense spectrums:* Many solutions struggle to operate effectively in densely occupied spectrum environments. For instance, sparse Fourier transform algorithms [10, 13, 41] and modulated wideband converter (MWC)-based subsampling [27] assume sparsely occupied frequency bands (occupancy rates capped at 5% to 10%). Even deep learning-based approaches [28] operate only up to 0.6× the Landau rate. However, this assumption often does not hold in real-world scenarios [12], especially for DSS use cases. (2) *Lack of real-time sensing capabilities:* Sequential spectrum scanning [12, 29] and narrowband monitoring approaches [5, 36] fail to achieve continuous real-time sensing across all frequency bands. This limitation can result in missed detection of rapidly changing signals, such as radar waveforms, and introduce signal distortion due to fast scanning processes. (3) *High power consumption:* Many solutions require custom hardware designs that result in power consumption comparable to that of ADCs operating at the Nyquist sampling rate [2, 17, 31]. This high power demand makes such methods unsuitable for low-power or large-scale deployments. (4) *High production cost:* MEMS filter-based spectrum sensing methods [11, 15] are not only challenging to manufacture but also costly due to process variations and the complexity of customization, therefore less viable for widespread adoption.

In this paper, we present *RippleSense*, a novel approach leveraging Nyquist sub-sampling theory and frequency spectrum's natural folding property to achieve real-time wideband spectrum sensing with a single conventional low-speed ADC. At its core, RippleSense invents a novel approach named *Signature Ripple Injection*, which annotates the input signals of different frequency bands at the radio front-end so that the folded frequency spectrum can be reconstructed with post-processing as depicted in Figure 1. Specifically, the input signal is band-pass filtered to isolate each Nyquist zone, which is then labelled using a distinct sinusoidal voltage ripple (i.e., the signature ripple) produced by a modified gain stage control circuit. The signal is then combined and folded to baseband, where digital post-processing recovers the origin zones of the signal by detecting the distinct ripple frequencies.

With the above idea, we provided a hardware/software system co-design that achieves real-time spectrum sensing, covering up to 10 GHz bandwidth. To ensure the system's practicality, our design provided particular optimization on the following two aspects:

- Computational efficiency: The system integrates an efficient spectrogram reconstruction algorithm, which is further optimized for specific spectral characteristics, reducing computational overhead and ensuring real-time operation.
- Scalability: The system accommodates devices with varying processing capabilities - from the implementation on FPGA devices with GHz sampling rates to outlining the application on low-power devices with tens of kHz sampling rates.

RippleSense is prototyped with RFSoC4x2 [39], supporting 5 GSPS sampling and 10 GB/s data rate. A parallelized, GPU-accelerated post-processing algorithm delivers 13 GSPS baseband performance for real-time processing. Experimental results show that the system can achieve a time resolution of 10 nanoseconds accross a bandwidth of up to 10 GHz and 4x sub-sampling, demonstrating high spectrum sensing capabilities.

In summary, RippleSense entails these contributions:

- An approach of injecting voltage ripples as signatures to differentiate signals from different Nyquist zones, allowing scalable and efficient sub-Nyquist sampling recovering wideband and dense spectrum occupancy.
- A high-performance end-to-end workflow for optimized processing efficiency and performance scalability, supporting a wide range of bandwidth requirements and hardware configurations.
- An open-sourced hardware and software co-design[2] capable of real-time spectrum sensing across 10 GHz of bandwidth, with comprehensive experiments analyzing trade-offs in sampling rate and temporal-frequency resolution.

2 Theory and Working Principle

In this section, we first provide a primer on the sampling properties of sub-Nyquist sampling, including aliasing and baseband folding. Then, we present the key observation that underpins the system's core signal processing methodology.

Primer on Aliasing and Nyquist Folding.

When sinusoidal functions are sampled at a frequency f_s, the resulting samples can fit multiple sinusoids of different frequencies. Sinusoids given by $\sin(2\pi(f + Nf_s)t)$, where N is an integer, yield identical sets of samples, leading to a frequency spectrum with equally strong responses at all these frequencies. This ambiguity makes these frequencies aliases of each other (Figure 4A). Alias frequencies can be expressed as positive values: $f_N(f) = |f + Nf_s|$. As f increases, the alias $f_{-1}(f)$ decreases, meeting at $f = f_s/2$, the Nyquist frequency (Figure 4B). The concept of folding is illustrated by the symmetry in the frequency spectrum around the Nyquist frequency (Figure 4C/D). Adjusting the frequency and amplitude of the sinusoid within a certain range shows this symmetry between 0 and f_s, known as folding. Based on the understanding of aliasing and

[2] https://github.com/andreaskuster/RippleSense

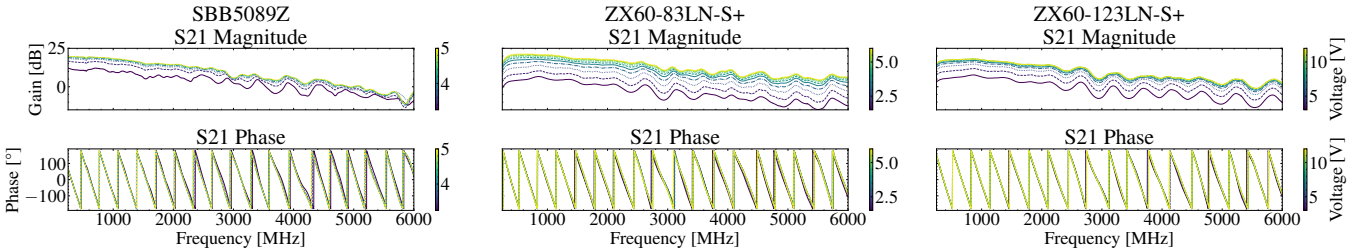


Figure 3: Forward transmission (S21) characterization of fixed-voltage RF amplifiers and their voltage-dependent gain and phase from 0.25-6 GHz.

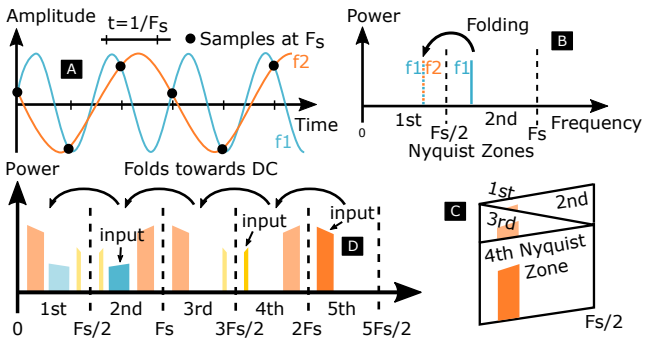


Figure 4: Nyquist sampling, aliasing and folding.

Nyquist folding theories, this paper proposes an innovative system design that can recover the original frequencies of all signals from the folded spectrum, thereby enabling full-band spectrum sensing.

Key Observation. The working principle of RippleSense is based on a key observation: *When a controllable time-varying change is applied to the input voltage of the RF front-end amplifier, this time-varying characteristic is preserved and reflected in the received signal.* Specifically, by adding a voltage ripple of a predefined frequency to the power supply of each fixed-gain amplifier, it modulates an amplitude modulation-like signal of that frequency on the incoming signal. Figure 3 demonstrates the variation of signal voltage and phase after passing through an amplifier supplied with variable voltage. Our characterization shows that all devices under test have a clear positive voltage/gain correlation and minimal phase shift. Note that while simple single transistor or opamp designs show a direct relation between supply voltage and gain, modern fixed-voltage gain amplifiers (as we used in the experiment), although more complex, still preserve the same positive correlation, even under signal saturation and temperature drift [4, 26]. This observation motivates the idea of injecting distinct signatures (i.e., the ripples) into the signals of each Nyquist zone before sampling at the lower baseband rate, so that we can preserve the information to identify the signal's origin after digitization without the necessity of complex additional hardware.

3 System Overview

Based on the key observation, we conduct a system design to recover the original frequencies of full band signals from the folded spectrum.

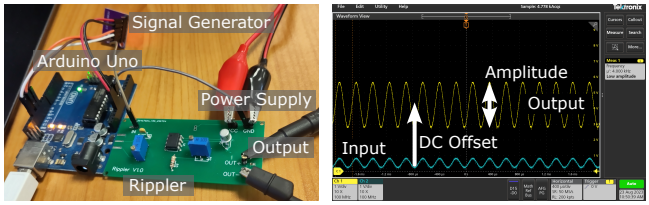


Figure 6: Supply voltage Rippler - physical setup with the custom PCB (tunable offset and amplitude) and oscilloscope analysis.

Figure 5 depicts the overall processing flow of the system, which is explained in this section.

Injection. After receiving the input signal (Figure 5A, one or multiple separate antennas) and filtering it via bandpass filters to separate each Nyquist zone's signals (Figure 5B), distinct signatures are inserted for distinguishability during reconstruction. Specifically, the signature is provided by a "Rippler" device, which supplies each gain amplifier with a fixed but configurable sinusoidal ripple generated by a programmable waveform generator (AD9833). Due to its limited voltage range and output power, we designed a custom PCB (Figure 6) that adds a configurable DC offset and ripple signal amplitude, followed by a BJT driver (2N2222A) to match the power demand of the gain amplifier. To reduce false positives during reconstruction, we spaced the signature ripple frequencies evenly over the temporal FFT spectrum (Alg.4:L4) keeping a gap from DC and from $\frac{F_s}{2}$ to prevent degradation due to DC offset removal respectively aliasing (Figure 15e). A sine wave as signature waveform was chosen for its efficient detection using FFT.

Mathematically we have

$$s(t) = A_c \cdot m_i(t) \cdot \cos(2\pi f_c t) \tag{1}$$

where $s(t)$ is the modulated signal, $m_i(t)$ is our i-th distinct sinusoidal ripple signature induced via a voltage ripple on the gain amplifier supply voltage. f_c is the carrier signal frequency and A_c is the carrier amplitude.

Detection and Reconstruction. The ripple detection and reconstruction algorithm is software-based and occurs after digitizing the analog signal. At this stage, signals from all higher Nyquist zones are combined (Figure 5D) in the baseband zone (Figure 5E), each with a distinct varying-gain-over-time signature, which appears over the time axis in each column. Computing the FFT of

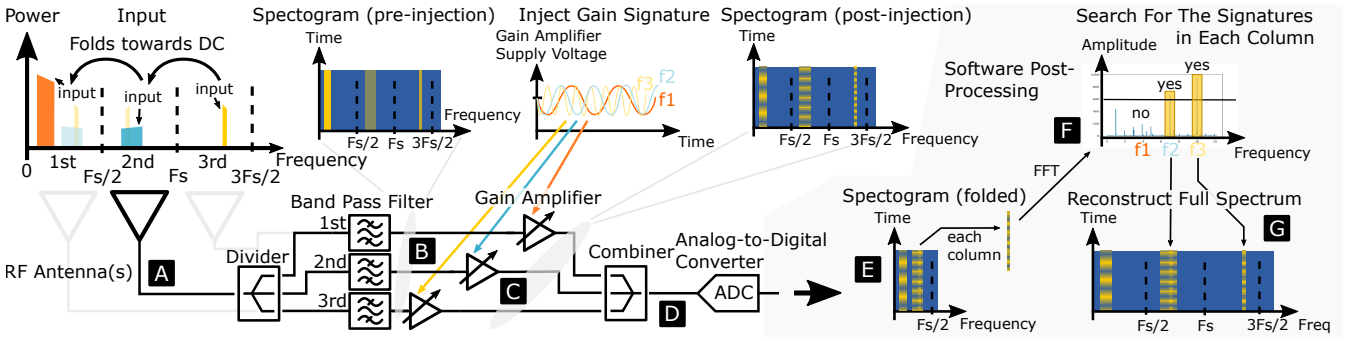


Figure 5: End-to-end hardware and software co-design from signal reception, splitting into the Nyquist zones, signature injection, and combination of the signal and sampling, followed by the reconstruction algorithm.

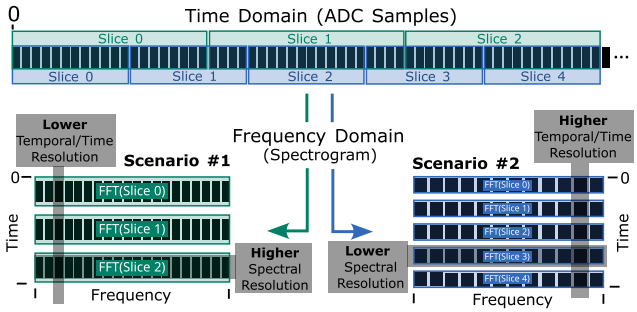


Figure 7: Illustration of the trade-off between temporal and spectral resolution in spectrum analysis.

the column extracts the frequency components present over time (Figure 5F). If there is a signal in the i-th Nyquist zone, it folds down to the baseband with the i-th ripple signature $m_i(t)$, present in the temporal FFT. Mathematically:

$$bb(t) = \sum_i A_i \cdot m_i(t) \cdot \cos(2\pi f_{bb} t) \tag{2}$$

where $bb(t)$ is the folded baseband signal and f_{bb} is the baseband alias of the higher Nyquist zones' carrier frequencies. If the i-th Nyquist zone has no active signal, A_i is zero, and $m_i(t)$ is absent in $bb(t)$ (Figure 5F). Using this information, we copy the column into all Nyquist zones with their signature present in the temporal FFT over $bb(t)$, reconstructing the entire spectrum (Figure 5G). The detailed reconstruction algorithm is presented in Section 4.

4 Spectrogram Reconstruction Algorithm

Post-processing is crucial for reconstructing the compressed, baseband-folded spectrum to its original form. This section examines each step and evaluates the overall performance.

Waterfall. The first step in the post-processing routine is to create a frequency spectrogram, or waterfall plot, of the incoming samples (Figure 5E). In practice, the data size must be limited to a finite duration to compute the discrete FFT. Choosing an optimal FFT window size is crucial as it influences the quality of the result.

Algorithm 1: Spectrogram Reconstruction

Input: Input samples X, num Nyquist zones n, spectral FFT size *fftsize*, ripple freq to Nyquist zone mapping $map \leftarrow \{(freq_i, zone_i)\}$;
Output: Reconstructed spectrogram Z;
Y ← waterfall(X, fftsize);
Z ← zeros(Y.shape[0] * n, Y.shape[1]);
for col_id, col ∈ enumerate(Y.cols) **do**
  temporal_fft ← FFT(col);
  for *zone, freq* ∈ *map* **do**
    if temporal_fft *contains* freq **then**
      if *zone* mod 2 == 0 **then** /* *even case* */
        Z[(zone+1)*Y.shape[0] - col_id] ← col;
      else /* *odd case* */
        Z[zone*Y.shape[0] + col_id] ← col;

As illustrated in Figure 7, a longer FFT processing window (i.e., a larger slice) in Scenario #1 enables derivation of a more fine-grained frequency spectrum, resulting in higher spectral resolution. However, the extended processing window may obscure rapid time variations, thereby reducing temporal resolution. In contrast, a shorter FFT processing window in Scenario #2 allows for more slices to be stacked, enhancing time resolution at the cost of lower frequency resolution.

This tradeoff could be overcome by using a sliding-window FFT; however, it would increase compute cost (e.g., an FFT window of 2048, single-element sliding-window utilizes 2048× more compute and memory) making it less energy-efficient and less scalable for high-speed ADCs and real-time use.

Reconstruction. The gain amplifier's source voltage ripple signatures are sine waves, detected using FFT in the time axis (temporal FFT). The waterfall plot from the previous step is used, and FFT is applied to each column, revealing the frequency components in that specific frequency bin over time. If an RF signal input is present in the i-th Nyquist zone, it leaves a trace of the ripple frequency and can be detected. Unlike the previous step, we do not apply a window function. The time resolution trace is potentially short, and losing information is detrimental. As our main interest is detecting

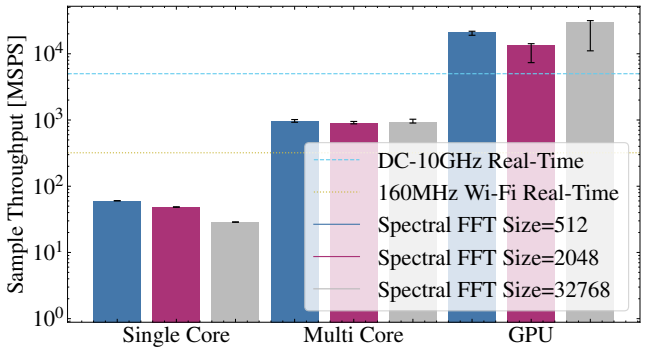


Figure 8: Algorithm performance on CPU and GPU.

pre-specified signature frequencies, spurious signals have limited implications. Even if signals from multiple zones perfectly fold onto each other in the sub-sampled spectrogram, we can recover all of them as long as the signatures do not alias each other, as shown in Section 6. Similarly, although perfect signal cancellation in baseband is theoretically possible, it is highly unlikely due to non-ideal responses from the antenna, ADC, and gain amplifier, plus our injected ripple. The maximum number of these overlaps determines the maximum Nyquist sub-sampling. With the ripple presence or absence information, we can reconstruct the entire spectrum by copying the columns into their original Nyquist zones (Figure 5G). Algorithm 1 provides an overview of these steps.

Computational Efficiency. Applications such as spectrum sharing require real-time signal processing, which is a core design goal. We evaluate the runtime complexity, and computational efficiency of our algorithm and its scaling from single-core to multi-core and GPU acceleration. The primary compute workload arises from the FFT operations of the spectrogram ripple detection and data copy/transfer. While highly optimized FFT libraries exist, we design and optimize each algorithm step to exploit parallelism and vectorization for maximum performance. Given N samples and an FFT window size W, the waterfall plot (Alg.1:L1) requires $\frac{N}{W}$ FFTs of length W, resulting in a complexity of $\frac{N}{W} * O(W \log W) = O(N)$. The FFTs in the time dimension (Alg.1:L4) involve W rows of length $\frac{N}{W}$, resulting in $O(N \log N)$, giving a total complexity of $O(N \log N)$. While theoretical runtime is important, practical implementation and optimization for real-time processing are required for an in-depth evaluation. Figure 8 shows the sample throughput of our implementation on our hardware platform[3] for different window sizes. The dashed lines indicate the minimum processing requirement for real-time sensing of 160 MHz (WiFi bandwidth) and 10 GHz spectrum bandwidth (our prototype system). With a 18.7x speedup on a 16-core/24-thread system, we achieve near-ideal parallelism on the CPU, further improved by a 270x speedup from single-core to GPU. With over 13 GSPS of GPU processing capacity, our implementation supports real-time processing even at peak sampling rates of the current ADC technology. Our single-core CPU implementation can process up to 48 MSPS, while the multi-core version achieves up to 904 MSPS. Memory requirements is

[3]Intel Core i9-12900KF processor, 64 GB DDR5-6000 memory (4x16 GB), Zotac GeForce RTX 3090 Ti, Python 3.10.12, Numpy 1.24.2, Pytorch 2.0.0+cu117, Cuda Toolkit 11.7

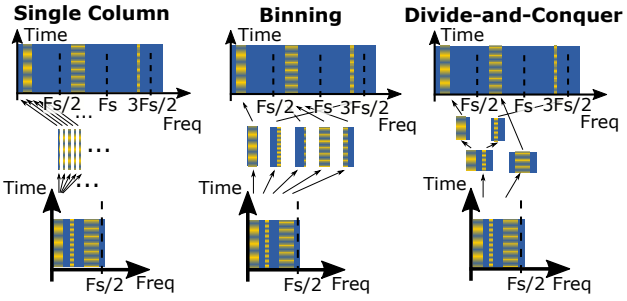


Figure 9: Performance optimization approaches.

scalable, and depends linearly on the batch processing size, which is important for parallel processors (multi-core and GPU). The limiting factor on GPU is not computation, given a large enough batch, but the PCIe 4.0 x16 bus with its theoretical 31.5GB/s bandwidth, where our 10GB/s data stream comes close. For that reason our GPU implementation makes use of large batches, to reduce data transfer overhead.

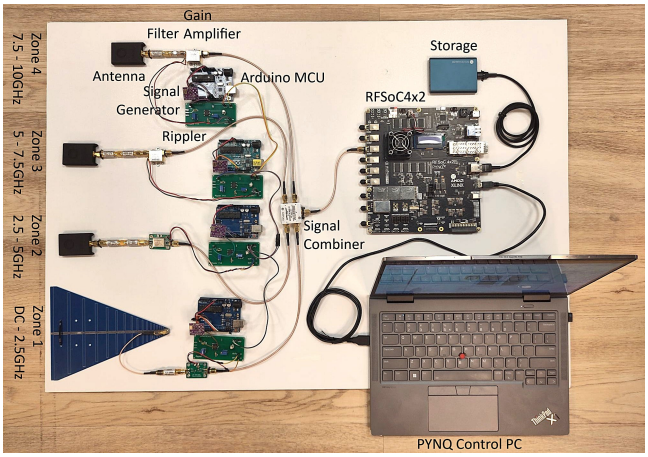
Single Line vs Bins vs Divide-and-Conquer. The presented algorithm is efficient and highly parallelizable, achieving 13 GSPS and delivering real-time performance on a single commodity GPU. However, specific optimizations can further increase throughput or reduce energy consumption and hardware demand (Figure 9). For many use cases, identifying the presence or absence of signals in ranges (e.g., 20, 40, 80, or 160 MHz wide Wi-Fi channels) is sufficient. After creating the initial spectrogram, we can sum the column values over the desired frequency bins and then apply the ripple detection and reconstruction algorithm to save computational resources (Table 1). Alternatively, we sample a few columns to derive the origin zone or sub-sample the input. For sparse spectra, a divide-and-conquer algorithm reduces computations. The spectrogram is split into two partitions, and the ripple detection algorithm is applied over the summed partitions. If at most one signature is present, the signal block's origin is determined, and the process terminates. If multiple signatures are present, the zone is further divided. Table 1 shows the operations comparison where sum is O(n) and FFT O(n*log(n)). The remainder of this paper focuses on the single-column variant, as it makes no assumptions about channel structure or signal sparsity.

Task	Single	Binning	Divide & Conquer
Waterfall	$\frac{N}{W}$ FFT(W)	$\frac{N}{W}$ FFT(W)	$\frac{N}{W}$ FFT(W)
Col Comb.	–	(W-B)SUM($\frac{N}{W}$)	$W \log(S)$SUM($\frac{N}{W}$)
Temp. FFT	W FFT($\frac{N}{W}$)	B FFT($\frac{N}{W}$)	$S \log(S)$FFT($\frac{N}{W}$)

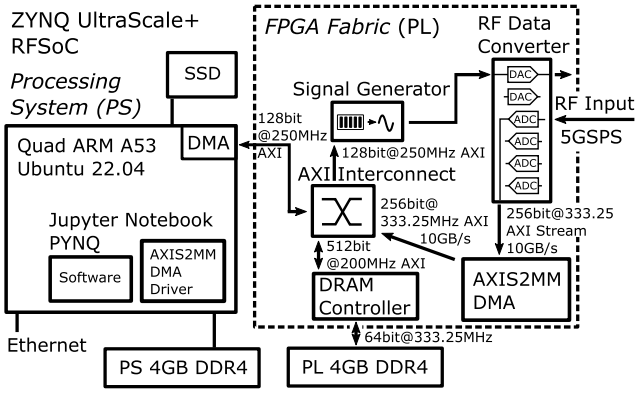
Table 1: Operation counts. N: samples, W: window size, B: binning factor, S: number of signals in baseband.

5 Scalability

This section focuses on the scalability optimizations of RippleSense, showcasing its ability to adapt to diverse performance and efficiency requirements. The section contains two parts. First, we discuss the system's capability to handle large volumes of data at GHz bandwidths, enabling high-performance scaling. Second, we outline



(a) End-to-end RFSoC4x2 RippleSense prototype.



(b) Functional overview of the FPGA design.

Figure 10: High-performance RF design enabling data processing and storage at a sampling rate of 5 GSPS.

strategies for energy and cost efficiency, facilitating efficient downscaling. Together, these approaches provide a versatile design space in terms of front-end design, computational power, and sampling rate, allowing the system to be tailored to specific application needs. Additionally, we explain how RippleSense can be integrated with compatible software-defined radios (SDRs) for rapid development.

5.1 Scaling Up

Applications that employ high-speed sampling FPGA devices for wideband spectrum analysis require the system to be capable of scaling up. In this case, the main challenge is managing the large data volume generated by data converter frontends, which may produce a continuous data streams of several gigabytes per second. For example, the RFSoC4x2 board uses a 14-bit / 5 GSPS data converter frontend which produces a continuous 10 GB/s data stream. This rate exceeds the write capabilities of most SSDs and approaches the 100 Gbit/s network interface limit. The actual scaling limitation, limiting our prototype to 4x sub-sampling and 10Ghz spectrum, however, is not data transport or signal processing but the constraints of the FPGA device. Specifically the PCB substrate, ADC analog bandwidth and the components used. The evaluation section will go a step further and discuss the theoretical scaling limits (Table 2).

Antenna. As indicated in light grey in Figure 5A, the antenna choice is flexible. To save physical space and cost, a single antenna can be used, but in our case, with the goal of sensing from DC to 10 GHz, to the best of our knowledge, no economically affordable antenna with a reasonable gain over the whole bandwidth exists. As the architecture allows the use of multiple antennas, we omit the signal divider and use a dedicated wide-band antenna for each Nyquist zone[4]. While this setup is not optimized for HF and VHF frequencies, multiple HF/VHF narrow-bandwidth antennas could be added to pick up the lower RF signals.

Filter, Gain Amplifier, RF Combiner. The high-performance front end is tuned to 5 GSPS, covering Nyquist zones with boundaries multiple of 2.5 GHz, which we filter using a combination of low-pass and high-pass filters[5]. For signal amplification and signature injection stage[6], the criterion is to cover the required frequency range, in addition to low noise and high gain. Depending on the choice of the operational amplifier, it is advisable to remove any energy buffering elements, such as external capacitors or inductors, from the voltage source, to allow the ripple signal to propagate efficiently to the gain response. We chose a resistive power combiner[7] due to its wide bandwidth down to DC. Figure 10a presents the end-to-end physical setup.

Signal Acquisition. The Xilinx RFSoC4x2 board [39] serves as the core component for GHz RF spectrum sampling, featuring two 9.85 GSPS transmitters, four 5 GSPS receivers, a reconfigurable FPGA, and four ARM Cortex-A53 processors. Two main challenges are addressed. First, RippleSense must meet hard real-time constraints for capturing, forwarding, processing, and storing samples to avoid data loss. Second, the RFSoC+ architecture [37] offers flexibility but poses challenges due to the complex direct sampling frontend, lengthy re-synthesis, missing IP features and limited FPGA debug capabilities. The data converter block is increased to the maximum sampling rate of 5 GSPS, requiring clock tree and PLL reconfiguration and aligning data clock and bus widths to support the required bandwidth. A dedicated AXI Stream path with 10 GB/s bandwidth is created from the ADC to the DRAM controller. A separate DMA controller transfers data from Programmable Logic (PL) memory to Processing System (PS) memory or SSD for post-processing. See Figure 10b for a system overview. Despite RFSoC's 5 GSPS ADC capability, the default Xilinx AXI DMA IP block [38] supports only up to 2^{26} bytes (64 MB) transfers, corresponding to a capture time of only 6.4 milliseconds at line rate, severely limiting real-time operation. Therefore, we implemented a custom AXI-Stream-to-AXI DMA block including a software driver from scratch for line-rate capture and storage (in DRAM), which is available on GitHub[8].

[4]Nyquist Zone 2-4: DEEPACE UWB-1P 2.4 GHz - 7 GHz [9], Zone 1: Non-Branded UWB antenna tuned for 800 MHz - 6 GHz
[5]DC-2.5GHz: VLF-2500+, 2.5GHz-5GHz: VLF-5000+ and VHF-2275+, 5GHz-7.5GHz: VLF-7200+, VHF-5050+, 7.5GHz+: VHF-7150+
[6]DC-2.5GHz: Generic 0.1-4000 MHz 20 dB Amplifier, 2.5GHz-5GHz: Generic 5-6000 MHz 20dB Amplifier, 5GHz-7.5GHz: Mini-Circuit ZX60-83LN-S+, 7.5GHz+: Mini-Circuit ZX60-123LN-S+
[7]Mini-Circuits ZFRSC-4-842-S+
[8]https://github.com/andreaskuster/RippleSense

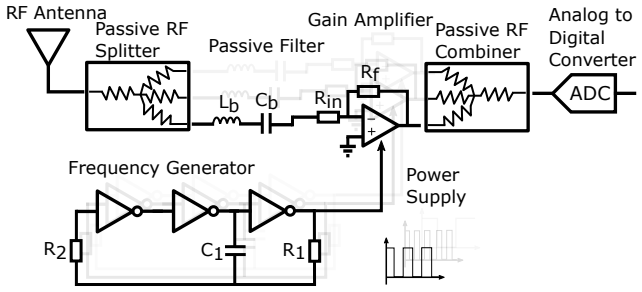


Figure 11: An alternative low-power, low-cost, and lower-performance RF frontend of RippleSense.

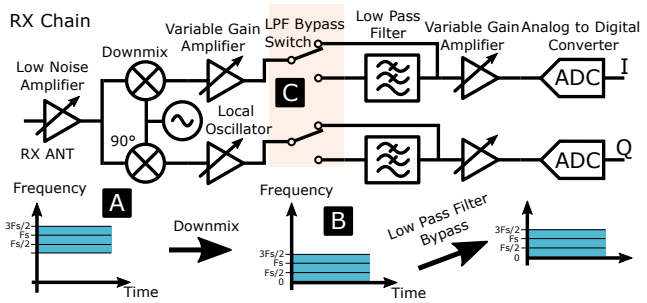


Figure 12: SDR-internal rx chain for RippleSense adoption. RX ANT input can be connected to the frontend in Figure 5 A-D.

This high-performance design has a peak power consumption of 250W: 5W front end, 40W RFSoC4x2, 80W CPU/System, and 120W GPU (average, under load). However, this represents the maximum power profile for our most demanding DC-10GHz sensing application; the system is designed to scale down to milliwatt levels for narrower bandwidth applications, as detailed in the next section.

5.2 Scaling Down

While we previously demonstrated that RippleSense can capture and process two-digit GHz of spectrum in real-time, an advantage is its parameter tradeoff capability and adaptability to different scenarios and applications. In this section, we outline what scaling down to the low-power wide area network (LPWAN) setting could look like. This paper, however, sets the primary focus on the large-scale prototype in order to be comparable to related works in Table 2.

Antenna. The antenna is a passive component and can be built or acquired without optimization required.

RF Divider and Combiner. The design considerations for signal divider and combiner stages focus on simplicity, two-way functionality, and passive component usage while ensuring RF input channel isolation. The device must be small, cheap, and easy to assemble. Dual use as both a divider and combiner simplifies design and reduces costs. The N-Way Wilkinson power divider (Fig11B/F) is ideal due to its simplicity, low cost, good isolation, and low signal loss. Its narrow frequency response and larger size at lower frequencies do not affect our LPWAN use case. If channel isolation is compromised, removing the resistors in the transmission path reduces signal loss and cost even further.

Filter. The bandpass filter design requires balancing simplicity, cost, and optimal signal response. A passive first-order Butterworth filter with a capacitor and inductor in series ($f = \frac{1}{2\pi\sqrt{L_b C_b}}$) meets these needs, reducing energy consumption (Fig11D). If needed, additional stages can be added to create a higher-order Butterworth filter.

Ripple Frequency Generator. The high-performance implementation used a signal generator with an MCU, consuming 193 mW (38.6 mA at 5 Volts) and costing a few dollars. However, RippleSense can tolerate less precise timings. A ring oscillator [22] or an inverter-based RC oscillator [16] can be used instead, costing only cents and microwatts. The ring oscillator frequency depends on gate delay and count, while the RC oscillator frequency is determined by resistor and capacitor values ($f = \frac{1}{2.2\pi R_1 C_1}$, for $R_2 = 10 * R_1$). This approach, shown in Fig. 11C, uses only three active components.

RF Gain Amplifier. To our knowledge, no passive, simple, and low-cost gain amplifier design exists. This critical stage injects the ripple signature, so an active, power-supplied component is needed for our already introduced injection method. An inverting operational amplifier (Fig. 11E) offers a simple design with configurable gain (gain $= -\frac{R_f}{R_{in}}$) and functions as a gain amplifier. Alternatively, a single-stage class A amplifier using one transistor can reduce component count and power consumption.

8-Channel LoRa Scenario. We modeled the circuit (Fig. 11) in Multisim for an application scenario involving observation of eight 125 kHz LoRa channels, focusing on energy and cost efficiency. For the ring oscillator, we use six 6-channel inverting buffers[9], and two 4-channel operational amplifiers[10] chosen for its suitable gain-bandwidth product. All other components are passive resistors, capacitors, and inductors. The total energy consumption of the active components is 108 µW for the buffers and 864 µW for the amplifiers, summing to 972 µW and a price point in the low single-digit dollars, which is within most LPWAN budgets.

Hardware Accelerated FFT. For low-power devices, FFT can be efficiently implemented as a hardware accelerator [14], providing high performance and low energy consumption.

Summary. This section showcases how the front-end design can be carefully designed to use cheap and simple commodity and mostly passive components, creating a low-power and low-cost front-end implementation. Alternatively, and for effective scaling to high subsampling factors, it can also be implemented in analog VLSI. Pairing this with a low-performance ADC and an efficient signal processing chain, it is viable to be included in commodity LPWAN applications to enhance its communication capabilities.

5.3 Compatibility with SDRs

Besides direct sampling RF devices such as the RFSoC4x2 [39], RippleSense can be implemented on compatible SDR platforms. The newest version of the RTL-SDR dongle supports direct sampling

[9] TI SN74LVC04ADR, 6*10 µA @ 1.8 V = 108 µW, USD 0.17
[10] TI TLV9001TIDCKR, 2*240 µA @ 1.8 V = 864 µW, USD 0.11

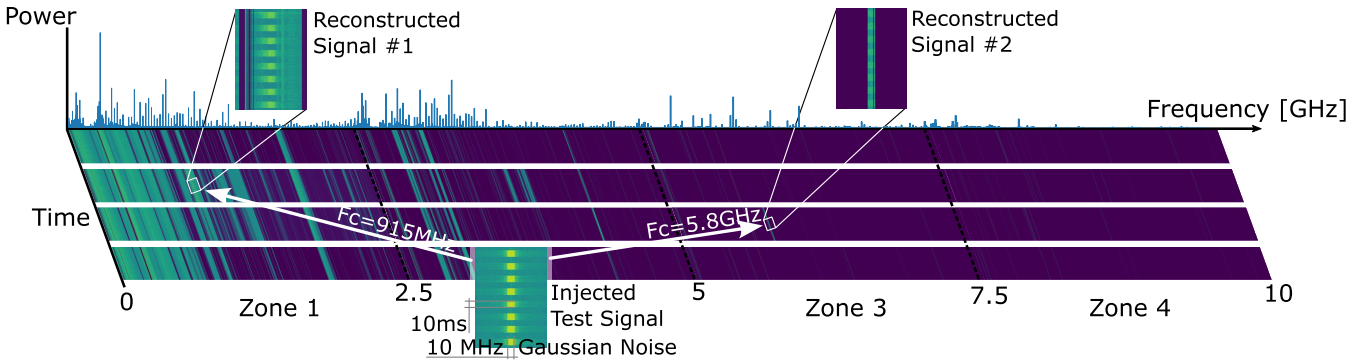


Figure 13: Spectrogram of the DC - 10 GHz real-world, large-scale field tests using the setup in Figure 10a. Four consecutive traces were recorded, with two traces containing reference signal injected at ISM frequencies

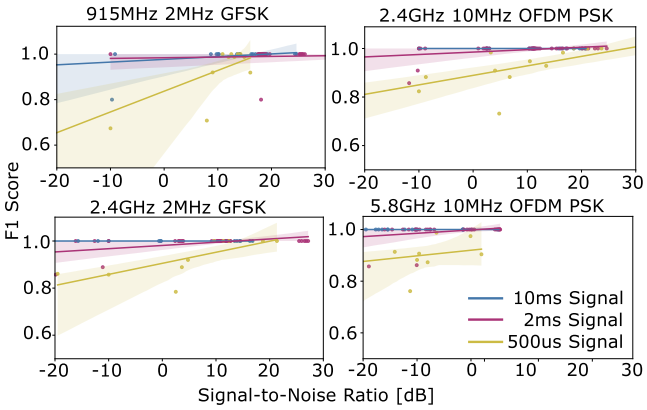


Figure 14: Over-the-air evaluation of our prototype 10a at ISM frequencies, bandwidths and modulations, at 95% confidence.

out-of-the-box [3], and popular devices such as the BladeRF [19] (LMS6002D [23]) and the LimeSDR [25](LMS7002D [24]) can be configured to add support, which we will elaborate next.

The input signal, consisting of bandwidths higher than the Nyquist sampling rate is received using an appropriate antenna (Figure 12A), and then downmixed to baseband (Figure 12B), followed by a signal amplification stage. The key lies in the low pass filter, originally configured to prevent aliasing which can be bypassed for BladeRF and LimeSDR devices as illustrated in Figure 12C. Together with the frontend (Fig.5A-C) connected to RX ANT, this is all that RippleSense needs, allowing integration with compatible SDRs.

6 Experiments

In this section, we evaluate the over-the-air performance of our DC–10 GHz setup (Figure 10a) and demonstrate how design parameters such as sampling rate, temporal and spectral resolution can be traded off to achieve the desired performance. The focus is on validating the RippleSense prototype and exploring parameter scalability, with a comprehensive evaluation of power aspects for scaling down left to future work.

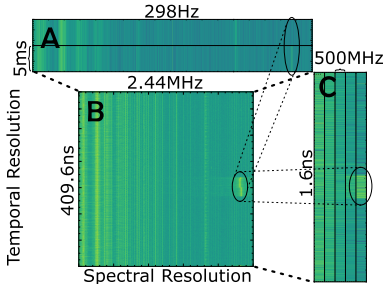
6.1 RFSo4x2 Over-the-Air Experiments

6.1.1 Test Bed. The test bed for our experiments consists of the RFSoC4x2 high-performance setup (Figure 10a) as receiver and a USRP B205mini [32] software-defined radio as transmitter injecting signals at 915 MHz, 2.4 GHz, and 5.8 GHz ISM bands.

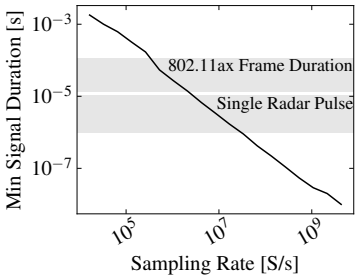
6.1.2 Full Spectrum Reconstruction and Signal Injection. The first experiment finalizes the end-to-end approach by reconstructing the DC–10 GHz RF spectrum over the air. To capture longer time durations than the DRAM buffer can hold, we stitched together four traces. We qualitatively verified the signal reconstruction by injecting a 10 MHz bandwidth RF test signal of Gaussian noise at 915 MHz and 5.8 GHz ISM bands, as shown in Figure 13. The signal has a 10 ms duration and a 20 ms period.

6.1.3 F1 Score and Signal-to-Noise Ratio. For quantitative performance analysis of different carrier frequencies, bandwidths, modulation techniques, signal durations, and signal-to-noise ratios, we transmitted and captured several thousand traces using the RippleSense prototype in four scenarios: 915 MHz GFSK for narrow-band technologies like ZigBee and LoRa, 2.4 GHz 2 MHz GFSK for Bluetooth, and OFDM PSK for 802.11 WiFi at 2.4 GHz and 5.8 GHz. To vary the SNR, we adjusted the transmit signal power and computed $\text{SNR} = \frac{P_{\text{signal+noise}} - P_{\text{noise}}}{P_{\text{noise}}} = \frac{P_{\text{signal}}}{P_{\text{noise}}}$, where $P_{\text{signal+noise}}$ is measured during transmission and P_{noise} during silent period. We report all values within our SDRs power limits, background noise, antenna performance, and amplifier gains.

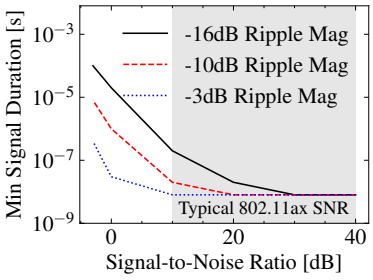
To reflect the dependency of false/true positives and negatives, we report the F1 score ($F1 = \frac{2 \cdot TP}{2 \cdot TP + FP + FN}$) as a point scatter plot, along with a regression model and 95% confidence interval. With higher SNR and longer signal traces, we expect the ripple signature to be captured more accurately due to processing gain, while carrier frequency and modulation effects should be minimal. Figure 14 shows that 10 ms signals are long enough for near-perfect performance in all scenarios. However, reducing the signal duration to 2 ms or 500 µs affects detection accuracy, especially at negative SNRs due to less processing gain. Some performance variation between cases can be attributed to non-flat antenna and amplifier gain responses, ADC non-linearity, cable loss and environmental factors (noise, temperature, ..). Despite this, performance remains stable



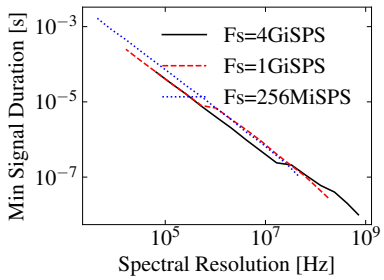
(a) Temporal vs. spectral resolution over-the-air (Fs=5GiSPS, duration=10ms).



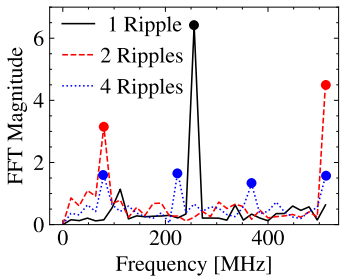
(b) Minimum signal duration vs. sampling rate (Fs=4GiSPS).



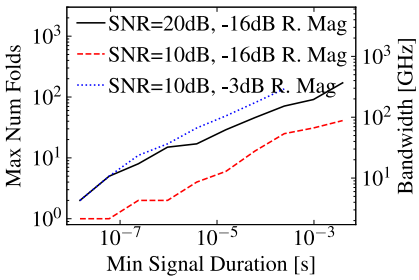
(c) Time resolution vs. SNR vs. ripple magnitude (Fs=4GiSPS).



(d) Minimum signal duration vs. spectral resolution (Ripple=-16dB).



(e) Influence of overlapping ripple signatures (Fs=1GiSPS, SNR=17dB, ripple=-16dB).



(f) Maximum Nyquist sub-sampling vs minimum duration (Fs=4GiSPS).

Figure 15: Parameter scalability and performance boundaries: (a) FFT resolution trade-offs, (b-d) minimum detectable signal duration vs. sampling rate, SNR, and spectral resolution, (e) signal overlap analysis, and (f) maximum supported Nyquist sub-sampling factors.

in all scenarios, achieving 70-90% accuracy down to -20 dB, which is well within the requirements for most RF protocol's operating range.

6.1.4 Temporal and Spectral FFT Resolution Tradeoff. In Section 4 we show the tradeoff between spectral and temporal resolution on a descriptive level, while we focus now on the effects in Figure 15a using real-world over-the-air samples. We build an intuition by decreasing the spectral resolution from 298 Hz (A) to 500 MHz (C) while keeping the number of samples (50 M) and the sampling rate (5 GSPS) constant, resulting in a capture time of 10 ms. As the spectral resolution decreases from A (298 Hz) to B (2.44 MHz) and C (500 MHz), the temporal resolution at A (5 ms) degrades significantly compared to B (409.6 ns) and C (1.6 ns). We analyzed a signal in B to interpret the limitations imposed on A and C. The signal in B occupies primarily the first half of the capture time. Due to the low temporal resolution in A, we can only distinguish between two time zones of 5 ms duration each. This prevents us from pinpointing the exact timing of the transmission. However, A provides a very high-frequency resolution, allowing us to observe the signal's spectral position clearly. In contrast, in C, we can precisely determine the signal's temporal position, but the frequency resolution is insufficient to reveal the slight frequency shifts over time observed in B. With a frequency resolution of 500 MHz in C, it becomes impossible to capture this variation.

This highlights the importance of selecting an appropriate trade-off between temporal and frequency resolution, depending on the application's requirements. This effect can be fully mitigated by either increasing the sampling rate, extending the capture duration or by using a sliding-window FFT, at the cost of a compute penalty.

6.2 Scalability Experiments

In the first part of our work, we focused on the evaluation of the RippleSense RFSoc4x2 prototype. However, to further evaluate the performance in terms of SNR, sampling rate and dense spectrum occupancy, some practical challenges arise. To our knowledge, no commercial device can sense or generate 10 GHz of bandwidth, preventing a complete ground-truth comparison for dense spectrum. In addition, much of the RF spectrum in this range is licensed, making it challenging to conduct over-the-air experiments (TX, in a noisy, non-Faraday cage setup) and to precisely control SNR, which requires transmit privileges. Therefore, for the second part of the experiment section, we address these challenges and showcase scalability and achievable limits by emulating the RF frontend from Figure 5A-D using Formula 1. Despite these constraints, our findings demonstrate consistency. A direct comparison between Figure 14 and Figure 15c (SNR vs. signal duration) reveals alignment in performance, reinforcing the validity of this combined approach. Unless stated otherwise, we use a 4 GiSPS sampling rate and -16 dB ripple magnitude.

6.2.1 Minimum Signal Duration and Sampling Rate. The first experiment evaluates the minimum signal duration that we can detect and correctly assign to the origin Nyquist zone. One of the shortest-lived signals relevant in practice are pulsed radars (short and long-range air search, maritime and weather) as low as 1 μs, which we aim as a lower bound. In Figure 15b, we plot the ADC sampling rate against the shortest detectable signal, with confidence (FFT peak) defined as a factor of two over the noise floor. At 4 GSPS, signals as short as 10 nanoseconds are detectable, exceeding the requirements for pulsed radar detection by three orders of magnitude. A sampling rate of 30 MSPS is sufficient to detect a single 1 μs radar pulse, and 10 MSPS can detect a single WiFi6 802.11ax frame in its shortest configuration (MCS11, 160MHz, CR 5/6).

6.2.2 Minimum Signal Duration and Signal-to-Noise Ratio. To understand how our approach scales with noise, we conducted experiments from below the noise floor (-3 dB) up to a signal-to-noise ratio of 40 dB for different ripple magnitudes, as shown in Figure 15c. The results indicate that while ripple magnitude significantly impacts lower SNR settings, all three curves eventually achieve 10 ns precision. Higher ripple magnitudes help reach this lower bound even at low SNR levels, relevant for typical WiFi signal strength. Notably, typical WiFi frame durations are in the microseconds to low milliseconds range, providing ample headroom and useful for trading off against other metrics. This shows that the design can be organically tuned to low SNR settings, by accumulating power over time.

6.2.3 Minimum Signal Duration and Bandwidth. Figure 15d examines the relationship between signal duration and spectral resolution, which should be reciprocal. Practical tests confirm this, although there are slight deviations, particularly at $F_s = 1$ GSPS. For instance, with a signal duration of approximately 1 μs, the spectral resolution is 4.186 MHz at $F_s = 4$ GSPS, but 8.260 MHz at $F_s = 1$ GSPS, which is slightly higher than 7.900 MHz at $F_s = 256$ MSPS, nevertheless confirming the theoretical dependency between all the scalability parameters in practice.

6.2.4 Number of Nyquist Folds and Minimum Signal Duration. This experiment showcases how many folds, respectively what subsampling factor can be supported by RippleSense (i.e. observable bandwidth = subsampling-factor $* \frac{F_S}{2}$), compared to the Nyquist sampling theorems limit. Note that only signals from different Nyquist zones, that are aliases of each other will overlap (see Sec.2). We first examine qualitatively how multiple overlapping signatures on the same baseband frequency (folded-down aliases) affect detection. As shown in Figure 15e, two effects influence the FFT peak. Overlapping ripples in the base band reduce each other's peak energy (SNR reduction, shown from black, red to blue), and the presence of more distinct ripple signatures increases the risk of false positives due to neighbouring peaks (frequency resolution/separation). Our experiments indicate that the second effect is less prominent, as the peak magnitude reduces quickly towards the noise level, leaving sufficient FFT bins between signal peaks.

Figure 15f shows that the number of distinguishable Nyquist zones scales with temporal resolution up to 60 folds at spectral FFT size of 2048. Achieving 60 folds would require significant front-end scaling, which might be feasible with analog integrated circuit designs. Practical applications, however, require fewer folds, like our high-performance prototype with four folds. With an SNR of 20 dB, scaling remains efficient, but at 10 dB, an FFT size of 256 is needed to distinguish a single fold. Increasing the ripple magnitude to -3 dB allows scaling to 10 folds with an FFT size of 4096, as higher ripple magnitudes are easier to distinguish from noise. Each data point was acquired from 100 repetitions, with each ripple present with a probability of 0.5, and no false positives or negatives reported.

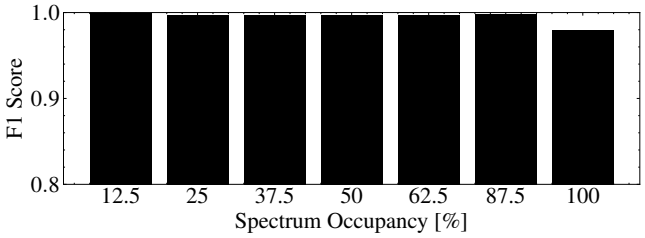


Figure 16: F1 score vs spectrum occupancy.

6.2.5 F1 Score and Spectrum Utilization. Next, we validate the performance on the densely occupied spectrum. We scaled from 0 to 100% spectrum utilization in Figure 16 using a sampling rate of $F_s = 256MSPS$, a signal-to-noise ratio of $10dB$, with a signal duration of 2 ms. We can observe that for low utilization, the F1 score is above 99% and slightly drops to 98% due to signal leakage to neighbouring FFT bins and, thus, false positives, which demonstrates consistent performance up to fully dense occupied spectrum.

Summary. Our experiments demonstrate the trade-offs and flexibility of our proposed spectrum sensing approach. Parameters like sampling rate, minimum detectable signal duration, spectral resolution, temporal resolution, ripple magnitude, and signal-to-noise ratio can be traded off to meet specific requirements. For example, if only long-lived signals are of interest, the sampling rate can be significantly reduced, saving costs and energy, or enabling the system to operate below the noise floor with high spectral resolution or extensive sub-sampling across multiple Nyquist zones.

7 Related Work

We compare key metrics to other works which are summarized in Table 2. While spectrum sweeping technologies such as SweepSense [12] and Beyond Sensing [29] achieve the highest Nyquist subsampling factor, they lack the required temporal resolution to capture short-lived radar signals, potentially even missing longer-lived transmissions from WiFi. The performance on this aspect is 5 to 9 orders of magnitude worse than our work. For seamless integration into commodity systems without hardware overhead, the number of ADCs and antennas is typically limited to one or two, which is the case for RippleSense. However, several other approaches such as Bigband [13], D-BigBand [35], Beyond Sensing [34], WISE [20] SigChord [28] and ISLA [15] require more than two antenna and ADC, making it difficult to integrate in commodity systems. Other technologies, such as Bigband [13], D-BigBand [35] and Xampling [27], SigChord [28] build upon the assumption that the spectrum

	Temporal Resolution	Max Sub-sampling	Max Bandwidth	Power Consumption	Price Estimate	Max Sensitivity	Min ADCs	Min Ant[3]
Bigband [13]	0.1 - 1 ms	6x	900 MHz	<54 W + P^0	~$9900 + P^0	-100 dBm	2x3[1]	3
D-BigBand [35]	0.1 - 1 ms	6x	900 MHz	<54 W + P^0	~$9900 + P^0	-100 dBm	2x3[1]	3
HackRF One [1]	750 ms	300x	6 GHz	<2.5 W + P^0	~$350	–	2[1]	1
SweepSense [12]	5 ms	200x	5 GHz	<18 W	~$3300	-60 dBm	1	1
Beyond Sensing [34]	50 ms	40x	4 GHz	<36 W + P^0	~$6600 + P^0	-100 dBm	2	2
S^3 [11]	40-100 μs	8.5x	418 MHz	<5 W + P^0	–	< -87 dBm[5]	2	1
Wi-Spy Air [21]	1250 ms	–	700 MHz	~150 mW	~$1000	-95 dBm	1	1
WISE [20]	0.5 - 2 ms	1x	900 MHz	< 100 W	< $200	-75 dBm	3	1
ISLA [15]	0.05-0.65 μs	16x	100 MHz	–	–	–	2	2
SigChord [28]	48 μs	5x[6]	800 MHz	~150-850 W	~$50-100k	-96 dBm[5]	5	1
RippleSense (range)	0.01 μs - 10s	1-60x	kHz - GHz	mW - 250 W	<$15 - $4000	>-100 dBm[5]	1	1
RippleSense (prototype)	2 ms	4x	10 GHz	~250 W	~$4000	-100 dBm[4,5]	1	4

Table 2: Comparison of state-of-the-art spectrum sensing methods. ([0]Software post-processing, [1]2 per SDR for I/Q, [3]Minimum antennas required, [4]Lowest SNR for F1 score of ≥ 0.95 for ≥ 2ms in Fig.14, [5]Ambient noise floor of -90 dBm (typically from -90 to -110 dBm), [6]Sub-sampling down to the Landau rate)

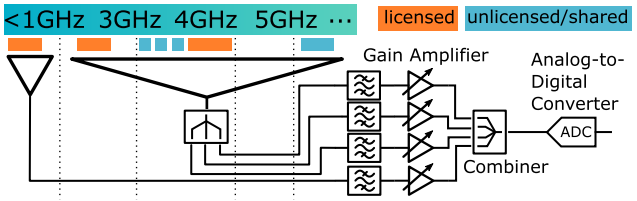


Figure 17: Handling disjoint 5G spectrum segments.

is sparsely allocated. While this might still hold true in some situations, it is certainly not a future-proof assumption, as part of our motivation is to share and re-use densely allocated spectrum, which is demonstrated in Figure 16. While MEMS resonator-based systems like ISLA [15] have gained traction in recent years and offer attractive performance, a custom filter needs to be designed for specific frequency and bandwidth requirements and cannot be bought off the shelf. Furthermore, according to [11], bandwidths below 20 MHz and above GHz cannot be sensed using a single filter. Other techniques such as WISE [20] based on UWB receivers make use of the higher bandwidth of UWB receivers but lack the scalability for Nyquist sub-sampling, which can only be circumvented by adding more receivers or waiting for technology scaling to higher bandwidths. To the best of our knowledge, our approach is the only one that can observe several GHz of the spectrum while having a temporal resolution down to two-digit nanoseconds, allowing us to observe short-lived signals emitted by radar.

8 Discussion

Non-Contiguous RF spectrum. Several wireless technologies such as Wi-Fi and 5G [30] rely on multiple communication channels or bands that are GHz apart. An advantage of our approach is its native support of such setups without performance penalties or processing overhead. The bands of interest can be filtered and folded down to the baseband for processing as shown in Figure 17.

Application: Channel Selection and CSMA for Commodity Devices. Ubiquitous communication technologies such as Wi-Fi and LoRa are deployed with the option of manually selecting the communication channel [33] to reduce collisions and allow higher capacity, reducing interference with neighboring deployments. Our approach enables opportunistic switching to the least-used channel by leveraging a holistic view of spectrum occupancy, and is scalable from high-bandwidth systems down to LPWAN/IoT.

Higher Bandwidth for Sensing Applications. Recent advances in sensing applications such as gesture recognition [18] and object detection performance benefit from higher RF bandwidth. Our approach can provide such an increase while keeping the power consumption and price low.

Spectrogram Learning Neural Network. The signal processing algorithm was kept efficient through high parallelizability and operation with low compute complexity to keep our implementation scalable to GHz and energy-efficient for low-power applications. If the compute resources can be increased for a specific application, an algorithm based on the detection of the ripple signature through spectrogram learning neural networks [40] can be applied, possibly scaling to lower signal-to-noise ratio or shorter signal durations.

Ripple Signature. The gain amplifier's supply ripple was deliberately designed using distinct sine wave frequencies for efficient FFT detection. Future work could explore more complex signatures with correlation or neural network detection, potentially simplifying the frontend.

9 Conclusion

This paper presents RippleSense, a novel hardware-software co-design that redefines the limits of wideband spectrum sensing through the injection of signature ripples. By leveraging the natural folding properties of sub-Nyquist sampling and annotating frequency bands at the radio front-end, RippleSense achieves real-time reconstruction of dense RF spectrum across bandwidths up to 10 GHz. Our evaluation demonstrates that this approach is scalable, from high-performance FPGA-based platforms with 10 ns temporal resolution and GHz of bandwidth to ultra-low-power LPWAN

applications operating at milliwatt levels. By outperforming state-of-the-art solutions in both bandwidth and temporal resolution, RippleSense provides a practical foundation for the next generation of dynamic spectrum sharing and real-time network monitoring. We believe this signature-based reconstruction paradigm opens new avenues for ubiquitous, high-bandwidth sensing in increasingly congested RF environments.

Acknowledgments

We would like to thank Prof. Rui Tan, Dr. Qirui Liu, and Dr. Songfan Li for their insightful discussions and help. We also thank our anonymous shepherd and reviewers for their constructive comments that helped improve this work. This work is supported by the Hong Kong Research Grants Council (RGC) General Research Fund (GRF) 16204224.

References

[1] [n. d.]. HackRF One Spectrum Sweep Mode. ([n. d.]). https://github.com/greatscottgadgets/hackrf/releases/tag/v2017.02.1.

[2] Omid Abari, Fabian Lim, Fred Chen, and Vladimir Stojanović. 2013. Why Analog-to-Information Converters Suffer in High-Bandwidth Sparse Signal Applications. *IEEE Transactions on Circuits and Systems I: Regular Papers* 60, 9 (Sep. 2013), 2273–2284. https://doi.org/10.1109/TCSI.2013.2246212

[3] RTL-SDR.com Admin. 2013. RTL-SDR Direct Sampling Mode. (2013). Retrieved August 30, 2023 from https://www.rtl-sdr.com/rtl-sdr-direct-sampling-mode/

[4] Analog Devices. 2023. *ADL5531 20 MHz to 500 MHz IF Gain Block.* Analog Devices Inc., One Technology Way, P.O. Box 9106, Norwood, MA 02062-9106. https://www.analog.com/en/products/adl5531.html Fixed-gain amplifier, 20 dB typical gain, 500 MHz bandwidth, 3.3 V to 5 V supply.

[5] A. A. Cheema and S. Salous. 2016. Digital FMCW for ultrawideband spectrum sensing. *Radio Science* 51, 8 (2016), 1413–1420. https://doi.org/10.1002/2016RS006083

[6] Federal Communications Commission. [n. d.]. White Space. ([n. d.]). Retrieved September 01, 2023 from https://www.fcc.gov/general/white-space

[7] Federal Communications Commission. 2015. 3.5 GHz Band Overview. (2015). Retrieved September 01, 2023 from https://www.fcc.gov/wireless/bureau-divisions/mobility-division/35-ghz-band/35-ghz-band-overview

[8] Federal Communications Commission. 2020. FCC Opens 6 GHz Band to Wi-Fi and Other Unlicensed Uses. (2020). Retrieved September 01, 2023 from https://www.fcc.gov/document/fcc-opens-6-ghz-band-wi-fi-and-other-unlicensed-uses

[9] DEEPACE. [n. d.]. DEEPACE Wideband Omnidirectional Antenna UWB+1S 2.4GHz-7GHz. ([n. d.]). Retrieved November 26, 2023 from https://deepace.net/product/uwb1s-2-4ghz-7ghz-wideband-omnidirectional-antenna/

[10] Badih Ghazi, Haitham Hassanieh, Piotr Indyk, Dina Katabi, Eric Price, and Lixin Shi. 2013. Sample-optimal average-case sparse Fourier Transform in two dimensions. In *2013 51st Annual Allerton Conference on Communication, Control, and Computing (Allerton).* 1258–1265. https://doi.org/10.1109/Allerton.2013.6736670

[11] Junfeng Guan, Jitian Zhang, Ruochen Lu, Hyungjoo Seo, Jin Zhou, Songbin Gong, and Haitham Hassanieh. 2021. Efficient Wideband Spectrum Sensing Using MEMS Acoustic Resonators. In *18th USENIX Symposium on Networked Systems Design and Implementation (NSDI 21).* USENIX Association, 809–825. https://www.usenix.org/conference/nsdi21/presentation/guan

[12] Yeswanth Guddeti, Raghav Subbaraman, Moein Khazraee, Aaron Schulman, and Dinesh Bharadia. 2019. SweepSense: Sensing 5 GHz in 5 Milliseconds with Low-cost Radios. In *16th USENIX Symposium on Networked Systems Design and Implementation (NSDI 19).* USENIX Association, Boston, MA, 317–330. https://www.usenix.org/conference/nsdi19/presentation/guddeti

[13] Haitham Hassanieh, Lixin Shi, Omid Abari, Ezzeldin Hamed, and Dina Katabi. 2014. GHz-wide sensing and decoding using the sparse Fourier transform. In *IEEE INFOCOM 2014 - IEEE Conference on Computer Communications.* 2256–2264. https://doi.org/10.1109/INFOCOM.2014.6848169

[14] Dongsuk Jeon, Mingoo Seok, Chaitali Chakrabarti, David Blaauw, and Dennis Sylvester. 2012. A Super-Pipelined Energy Efficient Subthreshold 240 MS/s FFT Core in 65 nm CMOS. *IEEE Journal of Solid-State Circuits* 47, 1 (Jan 2012), 23–34. https://doi.org/10.1109/JSSC.2011.2169311

[15] Suraj Jog, Junfeng Guan, Sohrab Madani, Ruochen Lu, Songbin Gong, Deepak Vasisht, and Haitham Hassanieh. 2022. Enabling IoT Self-Localization Using Ambient 5G Signals. In *19th USENIX Symposium on Networked Systems Design and Implementation (NSDI 22).* USENIX Association, Renton, WA, 1011–1026. https://www.usenix.org/conference/nsdi22/presentation/jog

[16] Junsoo Ko and Minjae Lee. 2019. A 1.8 V 18.13 MHz Inverter-Based On-Chip RC Oscillator with Flicker Noise Suppression Using Logic Transition Voltage Feedback. *Electronics* 8 (11 2019), 1353. https://doi.org/10.3390/electronics8111353

[17] J. N. Laska, W. F. Bradley, T. W. Rondeau, K. E. Nolan, and B. Vigoda. 2011. Compressive sensing for dynamic spectrum access networks: Techniques and tradeoffs. In *2011 IEEE International Symposium on Dynamic Spectrum Access Networks (DySPAN).* 156–163. https://doi.org/10.1109/DYSPAN.2011.5936202

[18] Dong Li, Jialin Liu, Sunghoon Ivan Lee, and Jie Xiong. 2023. Room-Scale Hand Gesture Recognition Using Smart Speakers. In *Proceedings of the 20th ACM Conference on Embedded Networked Sensor Systems (SenSys '22).* Association for Computing Machinery, New York, NY, USA, 462–475. https://doi.org/10.1145/3560905.3568528

[19] Nuand LLC. 2023. bladeRF. (2023). Retrieved August 30, 2023 from https://www.nuand.com/bladerf-1/

[20] Zhicheng Luo, Qianyi Huang, Rui Wang, Hao Chen, Xiaofeng Tao, Guihai Chen, and Qian Zhang. 2023. WISE: Low-Cost Wide Band Spectrum Sensing Using UWB. In *Proceedings of the 20th ACM Conference on Embedded Networked Sensor Systems (SenSys '22).* Association for Computing Machinery, New York, NY, USA, 651–666. https://doi.org/10.1145/3560905.3568541

[21] Metageek. 2019. Wi-Spy Air Hardware Data Sheet. (2019). Retrieved November 26, 2023 from https://support.metageek.com/hc/en-us/articles/360016560613-Wi-Spy-Air-Hardware-Data-Sheet

[22] Vratislav Michal. 2012. On the low-power design, stability improvement and frequency estimation of the CMOS ring oscillator. In *Proceedings of 22nd International Conference Radioelektronika 2012.* 1–4.

[23] Lime Microsystems. 2012. LMS6002D Multi-band Multi-standard Transceiver with Integrated Dual DACs and ADCs. (2012). Retrieved August 30, 2023 from https://limemicro.com/app/uploads/2015/10/LMS6002Dr2-DataSheet-1.2r0.pdf

[24] Lime Microsystems. 2017. LMS7002M FPRF MIMO Transceiver IC With Integrated Microcontroller. (2017). Retrieved August 30, 2023 from https://limemicro.com/app/uploads/2017/07/LMS7002M-Data-Sheet-v3.1r00.pdf

[25] Lime Microsystems. 2020. LimeSDR. (2020). Retrieved August 30, 2023 from https://limemicro.com/products/boards/limesdr/

[26] Mini-Circuits. 2023. *ZX60-83LN-S+ Low Noise Amplifier Datasheet.* https://www.minicircuits.com/pdfs/ZX60-83LN-S+.pdf

[27] Moshe Mishali and Yonina C. Eldar. 2010. Xampling: Analog Data Compression. In *2010 Data Compression Conference.* 366–375. https://doi.org/10.1109/DCC.2010.39

[28] Jinbo Peng, Junwen Duan, Zheng Lin, Haoxuan Yuan, Yue Gao, and Zhe Chen. 2025. *SigChord: Sniffing Wide Non-sparse Multiband Signals for Terrestrial and Non-terrestrial Networks.* Association for Computing Machinery, New York, NY, USA, 28–41. https://doi.org/10.1145/3711875.3729159

[29] Damian Pfammatter, Domenico Giustiniano, and Vincent Lenders. 2015. A Software-Defined Sensor Architecture for Large-Scale Wideband Spectrum Monitoring. In *Proceedings of the 14th International Conference on Information Processing in Sensor Networks (IPSN '15).* Association for Computing Machinery, New York, NY, USA, 71–82. https://doi.org/10.1145/2737095.2737119

[30] Qualcomm. 2020. Global 5G spectrum. (2020). Retrieved November 26, 2023 from https://www.qualcomm.com/content/dam/qcomm-martech/dm-assets/documents/20200625_mipi_-_5g_spectrum_-_dean_brenner_4.pdf

[31] Moslem Rashidi, Kasra Haghighi, Ashkan Panahi, and Mats Viberg. 2011. A NLLS based sub-nyquist rate spectrum sensing for wideband cognitive radio. In *2011 IEEE International Symposium on Dynamic Spectrum Access Networks (DySPAN).* 545–551. https://doi.org/10.1109/DYSPAN.2011.5936245

[32] Ettus Research. [n. d.]. USRP B205mini-i. ([n. d.]). Retrieved November 26, 2023 from https://www.ettus.com/all-products/USRP-B205mini-i/

[33] Danielle Saliba, Rodrigue Imad, and Sebastien Houcke. 2017. Wifi channel selection based on load criteria. In *2017 20th International Symposium on Wireless Personal Multimedia Communications (WPMC).* 332–336. https://doi.org/10.1109/WPMC.2017.8301833

[34] Lixin Shi, Paramvir Bahl, and Dina Katabi. 2015. Beyond Sensing: Multi-GHz Realtime Spectrum Analytics. In *Proceedings of the 12th USENIX Conference on Networked Systems Design and Implementation (NSDI'15).* USENIX Association, USA, 159–172.

[35] Lixin Shi, Haitham Hassanieh, and Dina Katabi. 2014. D-BigBand: Sensing GHZ-Wide Non-Sparse Spectrum on Commodity Radios. In *Proceedings of the 6th Annual Workshop on Wireless of the Students, by the Students, for the Students (S3 '14).* Association for Computing Machinery, New York, NY, USA, 13–16. https://doi.org/10.1145/2645884.2645887

[36] Sriram Subramaniam, Hector Reyes, and Naima Kaabouch. 2015. Spectrum occupancy measurement: An autocorrelation based scanning technique using USRP. In *2015 IEEE 16th Annual Wireless and Microwave Technology Conference (WAMICON).* 1–5. https://doi.org/10.1109/WAMICON.2015.7120376

[37] AMD Xilinx. [n. d.]. Zynq UltraScale+ RFSoC RF Data Converter. ([n. d.]). Retrieved November 26, 2023 from https://www.xilinx.com/products/intellectual-property/rf-data-converter.html

[38] AMD Xilinx. 2023. AXI DMA v7.1 LogiCORE IP. (2023). Retrieved November 26, 2023 from https://docs.xilinx.com/r/en-US/pg021_axi_dma/AXI-DMA-v7.1-LogiCORE-IP-Product-Guide

[39] Xilinx/AMD. [n. d.]. RFSoC4x2 Overview. ([n. d.]). https://www.rfsoc-pynq.io/rfsoc_4x2_overview.html.

[40] Zheng Yang, Yi Zhang, Kun Qian, and Chenshu Wu. 2023. SLNet: A Spectrogram Learning Neural Network for Deep Wireless Sensing. In *20th USENIX Symposium on Networked Systems Design and Implementation (NSDI 23)*. USENIX Association, Boston, MA, 1221–1236. https://www.usenix.org/conference/nsdi23/presentation/yang-zheng

[41] Rabia Tugce Yazicigil, Tanbir Haque, Michael R. Whalen, Jeffrey Yuan, John Wright, and Peter R. Kinget. 2015. Wideband Rapid Interferer Detector Exploiting Compressed Sampling With a Quadrature Analog-to-Information Converter. *IEEE Journal of Solid-State Circuits* 50, 12 (Dec 2015), 3047–3064. https://doi.org/10.1109/JSSC.2015.2464708

[42] Juhwan Yoo, Christopher Turnes, Eric B. Nakamura, Chi K. Le, Stephen Becker, Emilio A. Sovero, Michael B. Wakin, Michael C. Grant, Justin Romberg, Azita Emami-Neyestanak, and Emmanuel Candes. 2012. A Compressed Sensing Parameter Extraction Platform for Radar Pulse Signal Acquisition. *IEEE Journal on Emerging and Selected Topics in Circuits and Systems* 2, 3 (Sep. 2012), 626–638. https://doi.org/10.1109/JETCAS.2012.2214634