usenix
THE ADVANCED
COMPUTING SYSTEMS
ASSOCIATION

Efficient Wideband Spectrum Sensing Using MEMS Acoustic Resonators

Junfeng Guan, Jitian Zhang, Ruochen Lu, Hyungjoo Seo, Jin Zhou, Songbin Gong, and Haitham Hassanieh, *University of Illinois at Urbana-Champaign*

https://www.usenix.org/conference/nsdi21/presentation/guan

This paper is included in the
Proceedings of the 18th USENIX Symposium on
Networked Systems Design and Implementation.
April 12–14, 2021

978-1-939133-21-2

Open access to the Proceedings of the
18th USENIX Symposium on Networked
Systems Design and Implementation
is sponsored by

NetApp

Efficient Wideband Spectrum Sensing Using MEMS Acoustic Resonators

Junfeng Guan, Jitian Zhang, Ruochen Lu, Hyungjoo Seo,
Jin Zhou, Songbin Gong, Haitham Hassanieh
Univeristy of Illinois at Urbana-Champaign

Abstract – This paper presents S^3, an efficient wideband spectrum sensing system that can detect the real-time occupancy of bands in large spectrum. S^3 samples the wireless spectrum below the Nyquist rate using cheap, commodity, low power analog-to-digital converters (ADCs). In contrast to existing sub-Nyquist sampling techniques, which can only work for sparsely occupied spectrum, S^3 can operate correctly even in dense spectrum. This makes it ideal for practical environments with dense spectrum occupancy, which is where spectrum sensing is most useful. To do so, S^3 leverages MEMS acoustic resonators that enable spike-train like filters in the RF frequency domain. These filters sparsify the spectrum while at the same time allow S^3 to monitor a small fraction of bandwidth in every band.

We introduce a new structured sparse recovery algorithm that enables S^3 to accurately detect the occupancy of multiple bands across a wide spectrum. We use our fabricated chip-scale MEMS spike-train filter to build a prototype of an S^3 spectrum sensor using low power off-the-shelf components. Results from a testbed of 19 radios show that S^3 can accurately detect the channel occupancies over a 418 MHz spectrum while sampling 8.5× below the Nyquist rate even if the spectrum is densely occupied.

1 Introduction

The past decade has witnessed significant changes in the wireless spectrum as the FCC (Federal Communications Committee) has repurposed many frequency bands for dynamic spectrum sharing. This includes the 6 GHz band, released in April 2020, to be shared between Wi-Fi 6E and the incumbent users in this band like microwave backhaul [14]. Another example is the 3.5 GHz Citizens Broadband Radio Service (CBRS) band, which was recently approved for commercial deployments in September 2019. To leverage the CBRS band, unlicensed devices must sense a 200 MHz spectrum and avoid causing interference to primary and licensed users like military radars [13]. Of course, an earlier and more well-known example of spectrum sharing is the TV White Spaces which were released in 2010 [15]. Moreover, there are lots of opportunities for spectrum sharing in the millimeter-wave frequencies. In particular, the FCC released 14 GHz of unlicensed spectrum in the 60 GHz band that can be shared among Wi-Fi and IoT technologies [16]. These changes have been driven by the ever-increasing demand for wireless connectivity and aim to exploit previously underutilized frequency bands to accommodate new unlicensed applications and achieve highly efficient usage of the spectrum.

Efficient and truly dynamic spectrum sharing, however, requires unlicensed devices to sense wideband spectrum (hundreds of MHz to GHz) in real-time to spot and access momentarily idle channels. Unfortunately, real-time wideband spectrum sensing is challenging since it requires high-speed analog-to-digital converters (ADCs) that can sample the signal at the Nyquist sampling rate. Such high-speed ADCs are expensive, have low bit resolution, and can consume several watts of power [4, 12, 19, 34, 58].[1] To avoid using high-speed ADCs, today's systems sequentially scan the spectrum, monitoring each narrow band for a short period of time [11, 47]. As a result, they cannot continuously sense all bands in real-time and can easily miss highly dynamic and fleeting signals such as radar waveforms in the CBRS band [53].

Past work has proposed using compressive sensing or sparse Fourier transforms to sense wideband spectrum without sampling at the Nyquist rate [21, 27, 31, 41]. However, these approaches inherently rely on the assumption that the frequency spectrum is sparsely occupied. Hence, they only work in the case of underutilized spectrum where at most 5% to 10% of the frequency bands are occupied [21, 58]. The goal of dynamic spectrum sharing, however, is to efficiently utilize the spectrum. Hence, wideband spectrum sensing must work even in a densely occupied spectrum in order to scale usage to many users and achieve high utilization.

In this paper, we introduce S^3 (Spectrum Sensing Spike-train), an efficient low power spectrum sensing system that can monitor the real-time occupancy of multiple frequency bands in a wide spectrum. S^3 samples the wireless spectrum below the Nyquist sampling rate using cheap, commodity, low power ADCs but does not assume that the spectrum is sparsely occupied. A key enabler of S^3 is the use of MEMS (mirco-electro-mechanical-system) acoustic resonators that can create a spike-train like filter in frequency as shown in Fig 1. The MEMS filter processes the signal in the acoustic domain using carefully designed piezoelectric resonators with an assortment of equally spaced resonance frequencies. The resonators will pass the signals in these resonance frequencies and filter out the rest before converting the signal back to the RF domain. This creates an RF filter with very narrow, sharp, and periodic passbands across a wideband spectrum.

The spike-train filter enables S^3 to sample the spectrum in

[1] In fact, the power consumption of spectrum sensors is dictated by the ADC sampling rate as shown in [58]. Hence, we can significantly improve the energy efficiency by reducing the sampling rate.

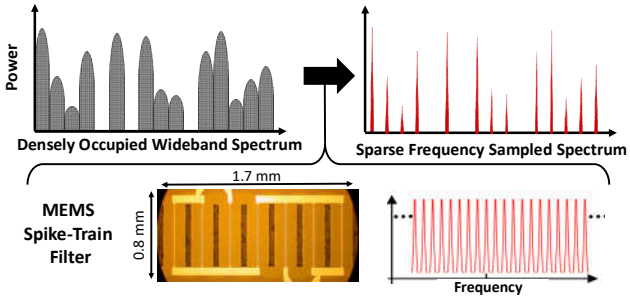

Figure 1: Filtering using MEMS acoustic resonators.

the frequency domain and monitor a small fraction of bandwidth in every band. S^3 can then tell if a band is occupied or idle by examining the sampled bandwidth in it, without the need to recover the entire band. This is like finding an available spot in a parking lot; We can tell if a spot is taken by peeking at some part of the car in it and we don't need to get close to every spot and see the entire car. Moreover, even if the wideband spectrum is densely occupied, the filter makes the spectrum significantly sparser as shown in Fig. 1. This enables S^3 to sample the signal below the Nyquist sampling rate and still recover the channel occupancies.

Translating S^3 into a practical system, however, requires addressing two key challenges. First, we need an algorithm that can accurately and efficiently reconstruct the spectrum occupancy. To address this, S^3 builds on past work in sparse recovery theory but differs from it in key aspects. In particular, compressive sensing algorithms require randomly sampling the time signal and cannot simply be implemented using low-speed ADCs [27, 59]. Sparse Fourier transform algorithms, on the other hand, can be implemented using low-speed ADCs, but they assume that the sparsely occupied bands are randomly distributed in the frequency spectrum [17, 21]. The MEMS filter creates a sparse spectrum that is highly periodic and far from random. For such sparsity patterns, sparse Fourier transform algorithms are highly sub-optimal.

S^3 aims to achieve the best of both worlds, i.e. no random sampling in time and no assumption of random distribution of occupied frequencies. To this end, S^3 leverages the uniquely structured sparsity pattern created by the filter to overcome the above challenges. The filter restricts the occupied frequencies to known locations in the spectrum, which significantly reduces the search space. It also allows us to optimize the sub-Nyquist sampling rate. In particular, optimal recovery can be achieved by choosing a sub-sampling factor that is co-prime to the number of spikes in the filter, as we show in section 5.

The second challenge is that in practice the MEMS resonators do not create an ideal spike-train. The spikes are not extremely narrow and have a small passband bandwidth which reduces the sparsity. Moreover, the separation between the spikes is not perfectly equal, and the spikes themselves are not identical. To address this, S^3 leverages the fact that different filters that are manufactured using the same process exhibit a very similar non-ideal spike-train, as we show in Sec. 6. Hence, the filter frequency response can be measured once and incorporated into the design of S^3. Specifically, we co-design the hardware and recovery algorithm of S^3 to account for the filter non-idealities and optimize its performance.

Evaluation: We had fabricated a chip-scale MEMS filter, shown in Fig. 1, which we leveraged to build a working prototype of S^3. The prototype can sense channel occupancies over a 418 MHz spectrum in real-time while sampling 8.5× below the Nyquist rate. The prototype uses two cheap, low power, off-the-shelf ADCs that sample around 50 MS/s ($\approx 1/17$ of the Nyquist rate). We extensively evaluate the performance of S^3 using a wireless testbed with 20 software defined radios that can occupy the entire 418 MHz spectrum at various power levels. Our results show that S^3 can accurately detect occupied channels. Even when the spectrum is as crowded as 90% occupied, S^3 achieves a false positive rate of 0.02 and a false negative rate of 0.0047. We also compare S^3 to state-of-the-art prior work like BigBand [21] and SweepSense [20] and demonstrate $5-10\times$ lower error rate for non-sparse spectrum. Furthermore, we show that S^3 can recover the wireless spectrum by performing outdoor and indoor measurements at various frequencies using a spectrum analyzer as the ground truth. Finally, we extend S^3 to not only detect the occupancy of the bands but also capture the power spectral density of the spectrum by quickly sweeping the center frequency for 22 MHz to cover the separation between the spikes.

Contributions: This paper has the following contributions:

- The paper bridges the latest advances in overtone MEMS acoustic resonators to RF spectrum sensing by leveraging spike-train filters to enable cheap and low power real-time wideband sensing of a densely occupied spectrum.
- The paper presents a novel sparse recovery algorithm that leverages the uniquely structured spectrum sparsity to efficiently recover a spectrum sampled significantly below the Nyquist sampling rate.
- The paper builds a prototype using commodity low-power components and evaluates its performance in a real testbed.

2 Background

In this section, we provide a brief background on wideband spectrum sensing using sub-Nyquist sampling. Further related work and background on spectrum sensing can be found in section 9.

This paper builds on past work that senses wideband spectrum without sampling at the Nyquist rate using compressive sensing [27, 31, 32, 39, 51, 58–60] or sparse Fourier transform algorithms [21, 43]. However, these approaches only work when the spectrum is underutilized and sparsely occupied which defeats the purpose of efficiently utilizing the spectrum. Furthermore, compressive sensing needs random

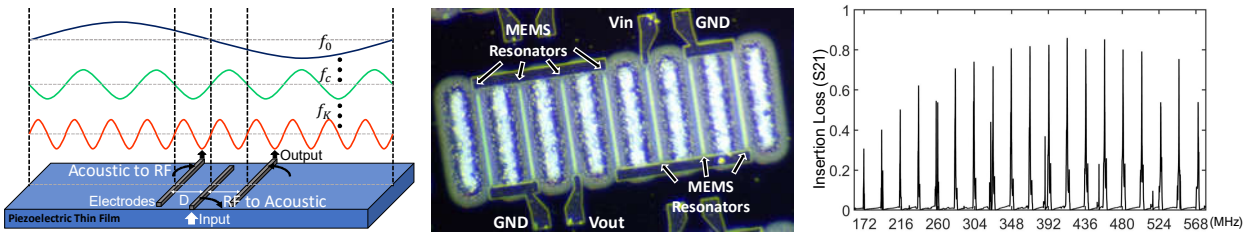

(a) Acoustic Resonator Diagram (b) Microscope Image of MEMS Filter (c) Spike-Train Filter Frequency Response

Figure 2: Spike-Train filter using MEMS acoustic resonators.

sampling [27, 58, 59], and as a result, requires custom hardware designs that can consume as much power as an ADC that samples at the Nyquist rate [1, 2]. Sparse Fourier transform algorithms do not necessarily require random sampling but must assume that the sparsely occupied bands are randomly distributed in the frequency spectrum to accurately recover the frequencies [17, 21].

BigBand [21] leverages sparse Fourier transform and uses co-prime sampling to acquire a sparse bandwidth while sampling 6× below the Nyquist rate. However, it only works up to 10% spectrum occupancy at which point it cannot recover the status of more than 14% of the spectrum. An extension, D-BigBand [43] can sense dense spectrum by considering the differential changes in occupancy. However, it assumes that the spectrum occupancy is mostly static with very few changes over time. Hence, it would not work for dynamic spectrum sharing where users sense and opportunistically transmit whenever they find an idle channel. [30] also attempts to extend BigBand to dense spectrum but requires sampling the signal first at the Nyquist rate in order to permute the samples and filter the signal before further sub-sampling it below Nyquist. S^3, on the other hand, can sense dense spectrum without the need for Nyquist sampling or random sampling. It also makes no assumptions on the changes in occupancy or the distribution of occupied bands across the spectrum.

3 Spike-Train MEMS Filter

Our work builds on recent advances in overtone MEMS RF filters [18, 40]. The MEMS filters convert RF signals into acoustic vibrations through the piezoelectric effect, then filter and process the signal in the acoustic domain before converting it back to the RF domain. Such filters can be further integrated with ICs to form an RF front-end solution, operating between a few MHz and 30 GHz for mobile and IoT devices. To this end, past work on MEMS RF filters optimize for a filter with a single passband [45, 64]. In contrast, the MEMS filter used by S^3 leverages overtone resonators that have an assortment of equally spaced resonance frequencies resulting in a spike train in the frequency domain. S^3 uses some of the very first spike-train MEMS filters which we had designed and fabricated [28, 29] to enable low power real-time wideband spectrum sensing of densely occupied spectrum.

To better understand how the MEMS filter works, consider the diagram of a MEMS acoustic resonator shown in Fig. 2(a). This resonator is commonly referred to as a LOBAR (Lateral Overtone Bulk Acoustic Resonator). The device consists of three electrodes on the top of a thin film made of the piezoelectric material $LiNbO_3$. RF signals come through the middle electrode and can be efficiently converted into acoustic waves through the piezoelectric effect, as long as their frequencies match the resonances of the film and are supported by the electrode design. Otherwise, the signals are reflected back and the frequencies are filtered out.

The resonance frequencies are determined by:[2]

(1) The width of the film: the film supports resonance frequencies for which acoustic wave vanishes at the edges of the film [8] i.e., the sine wave crosses zero at the edges as shown in Fig. 2(a). This condition is satisfied when the width of the film W is an integer (k) multiple of half a wavelength ($W = k\lambda/2$). Since $f = v/\lambda$, where v is the acoustic velocity in the piezoelectric material, the MEMS resonator will resonate at frequencies: $f_k = kv/2W$.

(2) The placement of electrodes: the filter will operate at center frequency f_c determined by the distance D between the electrodes: $f_c = v/2D$. Furthermore, for an odd number of electrodes, only acoustic waves that cross zero at the middle electrode, as shown in Fig. 2(a), will resonate.

Thus, the resonance frequencies will be the f_ks around f_c where k is even. This leads to a filter with center frequency f_c and a spike train where the spacing between the spikes is $\Delta f = v/W$. By modifying the width of the film and the position of the electrodes, we can modify Δf and f_c to control the frequency of the spikes in the filter.

The bandwidth or frequency span of the filter around f_c is determined by the electrodes where their RF-to-acoustic conversion efficiency degrades for resonance frequencies far from f_c, resulting in higher loss in spikes far from f_c. Adding more electrodes reduces the loss in spikes near f_c but narrows down the frequency span. We found that a three electrodes give the widest span with minimal loss of at most 2 dB.

Finally, to further enable a filter with very sharp and narrow

[2]The filter design actually involves a 4-way trade-off between (1) the frequency span, (2) the spacing between adjacent spikes, (3) the insertion loss in the spikes, and (4) the out-of-band rejection. For simplicity, we focus on the resonance frequencies design, and the trade-off between the frequency span and insertion loss. More details on the MEMS filter design trade-offs are discussed in [29].

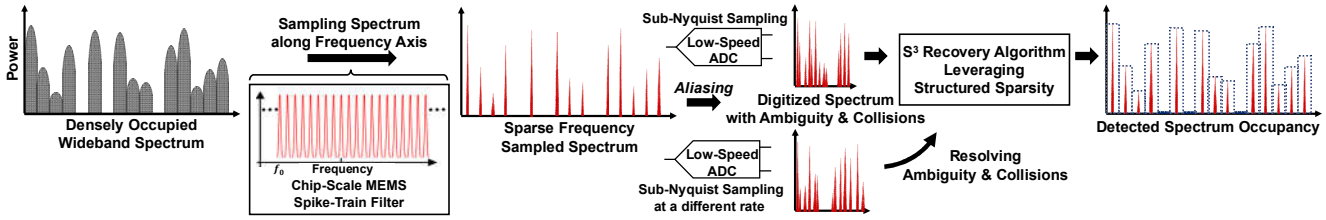

Figure 3: S^3 System Pipeline: S^3 samples the wideband spectrum along the frequency axis using the spike-train filter. The output spectrum is sub-sampled with low-speed ADC. S^3 leverages the structured sparsity in the filtered spectrum to resolve ambiguity and collisions due to aliasing and recovers spectrum occupancy.

spikes, we combine 7 of these MEMS resonators in a ladder filter topology [24]. Fig 2(b) shows an image of the filter under the Microscope. The frequency response of the filter is shown in Fig. 2(c). It has 19 periodic spikes with lowest loss between 161 and 579 MHz spanning a wide bandwidth of 418 MHz. The spacing between spikes is $\Delta f = 22$ MHz and the width of each spike is around $1 \sim 1.5$ MHz. Such an RF spike-train filter is the first of its kind. It presents a unique opportunity for processing wideband spectrum, which S^3 leverages towards efficient low power wideband spectrum sensing. It is also worth noting that the filter is passive and do not consume any power.

4 S^3 Overview

S^3 leverages the MEMS spike-train filters to sense wideband spectrum while sampling below the Nyquist rate. Fig. 3 illustrates an overview of the system pipeline. The received wideband spectrum is passed through the MEMS spike-train filter which samples the bands in the spectrum along the frequency axis. Specifically, the filter passes signals in frequencies aligned with the spikes and suppresses all the rest of the frequency components in the spectrum as shown in Fig. 3. The output of the filter is a sparse spectrum that preserves a small fraction of each band which we can use to monitor the occupancy of the band. Since the output spectrum is sparse, we can sample it below the Nyquist rate and still recover the occupancy information efficiently.

S^3 uses low-speed ADCs to sub-sample the signal. However, sampling below the Nyquist rate results in *"aliasing"* in the frequency domain i.e., multiple frequencies across the wide spectrum will alias (map) to the same frequency. Aliasing can lead to ambiguity and collisions, which prevent us from distinguishing frequencies that are occupied from those that are not. S^3 leverages the uniquely structured sparsity at the output of the spike-train filter to resolve such ambiguity and collisions and recover the spectrum occupancy. Ideally, one ADC is sufficient as we prove in section 5. However, due to practical limitations and imperfections in the spike-train filter, S^3 must use two ADCs sampling at different rates to accurately resolve ambiguity and collisions. We co-design the hardware and recovery algorithm to optimize the ADC sampling rates while accounting for the non-idealities of the spike-train filter as we describe in detail in section 6.

5 S^3 Recovery Algorithm

In this section, we describe S^3 recovery algorithm assuming an ideal spike-train filter. In later sections, we extend S^3 to deal with practical limitations.

Ideally, the spike-train filter will have equally spaced, very narrow and sharp spikes that can be approximated as an impulse train.[3] The frequency response of such a filter can be modeled as:

$$G(f) = \sum_{k}^{K} \delta(f - k\Delta f - f_0) \tag{1}$$

where K is the number of spikes, Δf is the spacing between spikes, and f_0 is the frequency of the first spike as shown in Fig. 3. Hence, the filter covers a spectrum bandwidth of $BW = \Delta f \times K$.

Let $x(t)$ be the input wideband signal in time domain and $X(f)$ be its non-sparse frequency representation whose bandwidth is also BW. After passing $x(t)$ through the spike-train filter, we get the signal $\tilde{x}(t)$ whose frequency spectrum is:

$$\tilde{X}(f) = X(f)G(f) = \sum_{k}^{K} A_k \delta(f - k\Delta f - f_0) \tag{2}$$

where $A_k = X(k\Delta f + f_0)$. $\tilde{X}(f)$ is at most K sparse i.e., it has at most K large frequency coefficients. Our goal is to recover these K coefficients A_k and estimate their power to detect the occupancy of the band around the frequency $f_0 + k\Delta f$.

S^3 samples the signal $\tilde{x}(t)$ using a low-speed ADC that samples at a rate $R = BW/P$ where P is an integer corresponding to the subsampling factor.[4] The sampling rate R is chosen such that $K \leq R \ll BW$. Let $y(t)$ be the sampled signal i.e., $y(t) = \tilde{x}(P \times t)$, and let $Y(f)$ be the Fourier transform of $y(t)$. Then, $Y(f)$ is an aliased version of $\tilde{X}(f)$:

$$Y(f) = \sum_{i=0}^{P-1} \tilde{X}(f + iR) \tag{3}$$

$Y(f)$ will cover a narrow bandwidth equal to R where frequencies in $\tilde{X}(f)$ that are equally spaced by R alias and sum together in the same frequency bin in $Y(f)$. Hence, once S^3 detects power in a frequency bin $Y(f)$, it knows that this power

[3] We can approximate the spikes as impulses if the width of the spike $\ll 1/T$ where T is the time window over which we sample the signal.

[4] Note that for simplicity, we have assumed that the ADC takes complex samples of the signal i.e., there are two ADCs sampling the I and Q of the wireless signal. We will relax this assumption in the following section.

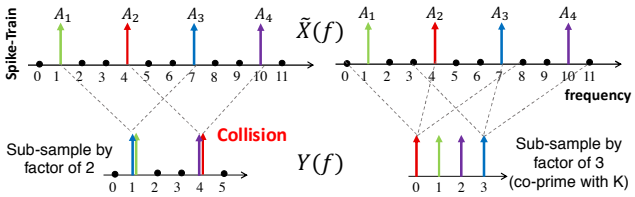

Figure 4: Co-prime sub-sampling factor avoids frequency collisions: Sub-sampling the spike train by 2 results in collisions between the spikes, but sub-sampling by 3 can avoid collisions because it is co-prime with K=4.

could have come from P different candidate frequencies in $\tilde{X}(f)$. Fig. 4 shows an example where if we sub-sample the signal by a factor $P = 2$, then every two equally spaced frequencies in $\tilde{X}(f)$ map to one value in $Y(f)$. Since $\tilde{X}(f)$ only has power in K coefficients A_k corresponding to the spikes of the filter, S^3 can easily eliminate a lot of candidates. Ideally, we want these coefficients to map to different bins. In this case, the bin value will be the same as the coefficient A_k which we can immediately estimate. However, if two coefficients A_{k_1} and A_{k_2} collide in the same bin as shown in Fig. 4, it will not be possible for S^3 to distinguish and estimate them.

S^3 can choose the sampling rate R in a manner that guarantees that no two coefficients collide. In particular, if the sub-sampling factor P and the number of spikes K are co-prime, then we can guarantee that none of K coefficients collide in the same bin and become indistinguishable. To see this, consider the example shown in Fig 4 where we have $K = 4$ spikes with coefficients A_1 to A_4 in the filtered spectrum $\tilde{X}(f)$. When we sub-sample by a factor of 2 below Nyquist, there will be collisions between A_1 and A_3, as well as A_2 and A_4. However, when we sub-sample by a factor of 3 below Nyquist, none of the coefficients collide, because the sub-sampling factor $P = 3$ and the number of spikes $K = 4$ are co-prime.

It is worth noting here that even though $P = 3$ uses a lower sampling rate than $P = 2$, increasing the sampling rate in this case results in more collisions. This is in contrast to past work on sub-Nyquist sampling [17,21] where higher sampling rates reduce collisions as the coefficients are assumed to be randomly distributed in the spectrum. Unlike past work, the structured sparsity of our spectrum requires carefully selecting the sampling rate to ensure that all coefficients can easily and immediately be recovered.

The below lemma, theoretically proves that if P and K are co-prime, then none of the coefficients will collide.

Lemma 5.1. *Given K, P are co-prime integers, let f_i and f_j be the frequencies of any two spikes in the spike train filter i.e. $f_i = k_i\Delta f + f_0$ and $f_j = k_j\Delta f + f_0$ such that $0 \le k_i, k_j < K$. Then, for all $f_i \neq f_j$, we have $f_i \neq f_j$ mod R.*

Proof. Assume there exist an $f_i \neq f_j$ such that the coefficients collide i.e., $f_i = f_j$ mod R. Note that by definition of the spike train, we also have $f_i = f_j$ mod Δf. Consequently, f_i and f_j are equal modulo the least common multiple: LCM$(R, \Delta f)$ = LCM$(BW/P, BW/K) = BW$, since K and P are co-prime.

Algorithm 1 S^3 Sensing with an Ideal Spike-Train Filter

Input: $x(t)$
$B_k \leftarrow$ Band around frequency $f = k\Delta f + f_0$
$\tilde{x}(t) = g(t) \circledast x(t)$ ▷ Filter $\tilde{X}(f) = X(f)G(f)$
$y(t) = \tilde{x}(P \times t)$ ▷ Sub-Nyquist Sample
$Y(f) = \text{FFT}(y(t))$
$A_k = Y\big((k\Delta f + f_0) \bmod R\big)$
if $E\left[|A_k|^2\right] > \sigma^2$ **then**
  B_k is occupied
else
  B_k is empty

Thus, $f_i = f_j$ mod BW which is a contradiction since BW is the entire bandwidth and we are given that $f_i \neq f_j$. Hence, by contradiction, for all $f_i \neq f_j$, we have $f_i \neq f_j$ mod R and none of the K coefficients collide. □

Given that we can choose a sampling rate that results in no collisions, we can easily recover the coefficients A_k as follows. We can compute $Y(f)$ by taking an FFT of $y(t)$ and for $0 \le k < K$, we directly set $A_k = Y((k\Delta f + f_0) \bmod R)$. We then apply an energy detector on A_k to obtain the occupancy of the band around the frequency $k\Delta f + f_0$. If $|A_k|^2$ is above the noise floor, then the band is occupied, otherwise, it is empty. A pseudocode for the overall sensing of S^3 with an ideal spike-train filter is shown in Alg. 1.

Next, we prove the below theorem about the correctness and the computational complexity of the algorithm.

Theorem 5.2. *Assuming a signal SNR > 0 dB for each occupied band, the system correctly recovers the occupancy of the bands using $O(K)$ samples and $O(K \log K)$ computations which is optimal.*

Proof. We will prove the above statement for the case where the entire spectrum is occupied. We can compute the the signal power of the filtered and sub-sampled signal as:

$$E\left[\|\tilde{Y}(f)\|_2^2\right] = E\left[\sum_{f=0}^{R-1} |Y(f)|^2\right] = E\left[\sum_{f=0}^{R-1}\sum_{i=0}^{P-1} |\tilde{X}(f+iR)|^2\right]$$
$$= E\left[\sum_{k=1}^{K} |\tilde{X}(k\Delta f + f_0)|^2\right] = E\left[\sum_{k=1}^{K} |A_k|^2\right] \ge KE\left[\min_k |A_k|^2\right] \tag{4}$$

Let σ^2 be the noise power per frequency. Since the spike-train filter suppresses the noise outside the spikes, the remaining noise in the signal is $K\sigma^2$. Hence, the SNR of the filtered and sub-sampled signal is:

$$SNR = \frac{E\left[\|\tilde{Y}(f)\|_2^2\right]}{K\sigma^2} \ge \frac{E\left[\min_k |A_k|^2\right]}{\sigma^2} > 1 \tag{5}$$

Thus, as long as the received signal is above the noise floor i.e. $SNR > 1$ (0 dB), filtering and sub-sampling will not increase the noise floor and the occupancy of the band can be detected

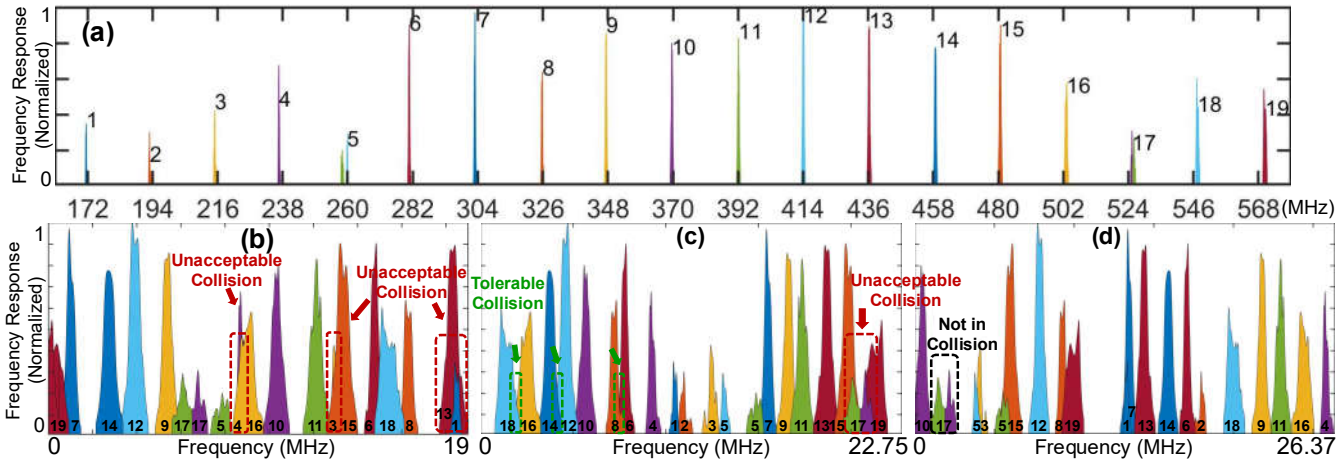

Figure 5: Aliasing of the spike-train filter at different sub-Nyquist sampling rates: (a) Locations of the 19 spikes on the frequency axis (b)Aliasing of the spikes at 38 MS/s (c) Aliasing of the spikes at 45.5 MS/s (b) Aliasing of the spikes at 52.74 MS/s.

correctly. Now, the algorithm samples at rate $R = O(K)$, takes an FFT of size $O(R)$ and then performs $O(R)$ computations. Hence, it requires $O(K \log K)$ computations and $O(K)$ sample, which is optimal. The algorithm is also deterministic, unlike compressive sensing and sparse Fourier transform algorithms which are randomized. □

6 S^3 with Practical Limitations

As mentioned earlier, the MEMS spike-train filter is non-ideal i.e., the spikes have some width as can be seen from Fig. 6. Although the ~ 1.5 MHz bandwidth is narrow compared to the channel bandwidth, it is still significant. Moreover, the spikes are neither identical nor perfectly equally spaced. In fact, they differ in magnitudes, bandwidths, and shapes. As a result, if we simply pick a sub-sampling factor P that is co-prime to the number of spikes K, there could be many collisions among the wide spikes. Figure 5 shows how the 19 spikes of our spike-train filter alias after sub-sampling. Figure 5(a) shows the spikes in the original wideband spectrum, while Fig. 5(b-d) show the aliasing of the spikes when sub-sampled at three different sampling rates. First, we choose the sampling rate to be 38 MS/s, because the resulting sub-sampling factor $P = 11$ is co-prime to $K = 19$. However, the aliased spectrum ends up with many collisions, as shown in Fig. 5(b). This suggests that the derived optimum no longer holds due to the practical limitations of the filter.

Fortunately, different filters that are manufactured through the same process exhibit a very similar spike train. Figure 6 compares the measured frequency responses of three spike-train filters we fabricated. We zoom into two spikes, otherwise the differences are very hard to spot. As one can see, the filters are almost identical. Hence, we can measure the frequency response of one spike-train filter and use it for the others.

Knowing the filter frequency response, we run an optimization problem to find a sampling rate that has as little collisions as possible. Ideally, this sampling rate should separate all the wide spikes after aliasing and prevent them from overlapping with one another. If a collision is unavoidable, we want it to only occur at the boundaries of the spikes, rather than having two wide spikes fully overlap. For example, as shown in Fig. 5(b), the collisions marked in red are unacceptable, because most of a spike's frequencies experience collision. In contrast, the collisions marked in green in Fig. 5(c) are tolerable, because only the boundaries of two spikes collide. Because we can simulate and compare the aliasing at different sampling rates offline, the optimization problem, in fact, can exhaustively search for all possible sampling rates.

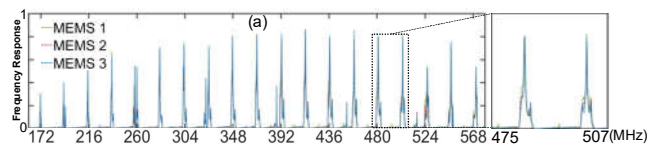

Figure 6: Measured frequency responses of three fabricated MEMS spike-train filters using the same process.

Another practical aspect is that in the real system, we only sample real signals and not complex in order to reduce complexity. Since the signal is real, the frequency representation is symmetric around the y-axis. Hence, with a sampling rate of $R = 38$ MS/s, the wideband spectrum actually aliases to a bandwidth of $R/2 = 19\ MHz$. Formally, if the original frequency of a spike is $f_{spike} = k\frac{R}{2} + b$, where $b < R/2$, then the aliasing frequency f_{alias} of the spike can be found through the following equation:

$$f_{alias} = \begin{cases} b & \text{if k is even} \\ \frac{R}{2} - b & \text{if k is odd} \end{cases} \quad (6)$$

In our specific case, we find 45.5 MS/s to be a really good sampling rate. As can be seen from Fig. 5(c), it spreads out the aliased frequencies of the wide spikes to 1.5, 2.6, 3.8, 4.6,

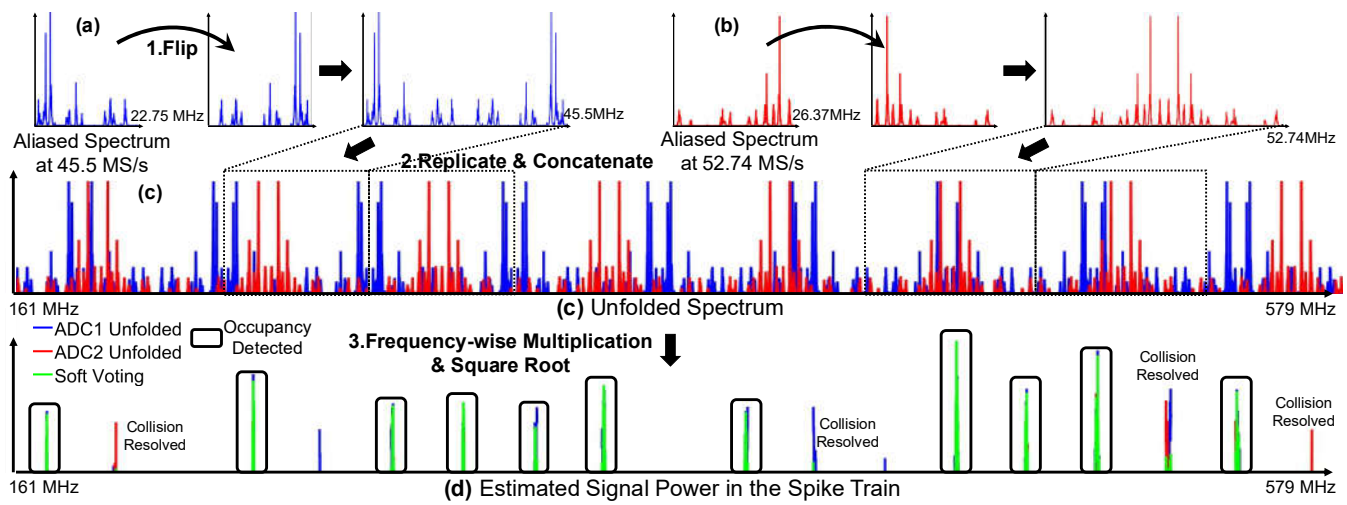

Figure 7: Practical S^3 recovery algorithm through voting: S^3 unfolds two aliased versions of the filtered spectrum to get a vote for the frequency components in the spike train. It then combines votes from the two different sampling rates to estimate the signal power in the spikes and the spectrum occupancy.

5.7, 7.3 MHz, etc. Therefore, most collisions only occur at the boundaries of the wide spikes. However, it still cannot avoid all unacceptable collisions. In fact, it's likely that no sampling rate can. For example, our spike-train filter has a unique 17th spike that is composed of two very close spikes. When sampling at 45.5 MS/s, these two small spikes completely overlap with the 15th and 19th spikes respectively. Therefore, when spike 15 and 19 are both occupied, we might falsely classify the spike 17 as occupied.

To resolve such unavoidable collisions, we leverage another sampling rate that provides us with a different set of aliasing frequencies for the spikes. We pick the second sampling rate in a way that any two spikes colliding at 45.5 MS/s do not collide again. To this end, we find a good sampling rate of 52.74 MS/s, and the resulting aliased frequencies of the spike train is shown in Fig. 5(d). One can see that the two parts of spike 17 do not collide with any other spikes at 52.74 MS/s. Thus, as long as we observe no power on frequencies corresponding to spike 17 at 52.74 MS/s, we will classify spike 17 as empty. Hence, by leveraging such incoherence between the two sampling rates, we can further resolve unavoidable frequency collisions and correctly identify the empty bands.

Using the two sub-Nyquist sampled spectra, S^3 recovers signal power in each spike, and then identifies the occupancy of the corresponding band. We leverage the two sampling rates through a soft voting scheme. The idea is that given an aliased spectrum and the sampling rate, we know all the possible original frequencies that correspond to the aliased frequencies. Hence, each aliased spectrum provides a vote for the source frequencies of the non-empty spectral components. Moreover, the non-empty frequencies on the original spectrum are also constrained to the spike-train frequencies. Therefore, when the two sampling rates vote for the same frequency that also falls in a spike, the frequency is very likely to be the true source frequency on the wideband spectrum.

Consider the two aliased versions shown in Fig. 7(a,b), where 11 out of the 19 bands are occupied and the other 8 bands are empty. Now we use them to vote where the non-empty frequency components come from. According to Eqn. 6, aliasing folds the wideband spectrum on to the bandwidth of $\frac{R}{2}$. Therefore, we can vote on all the possible source frequencies by unfolding the aliased spectrum. We accomplish this goal in the following three steps:

- **1. Unfold - Flip**: First, we flip the aliased spectrum $Y(f)$ with a bandwidth of $\frac{R}{2}$ to get the bandwidth between $\frac{R}{2}$ and R, as it equals to $Y(\frac{R}{2} - f)$ according to Eqn. 6.
- **2. Unfold - Replicate**: Then we replicate and concatenate the resulting spectrum from 0 Hz to R, and we get a vote for all frequencies in the frequency range of the spike train as shown in Fig. 7(c).
- **3. Soft Voting**: Finally, we combine the votes of the two sampling rates, where we only consider the frequencies within the spikes. This is done by multiplying the two votes on every frequency and taking a square-root. As a result, the non-empty frequencies that are voted by both sampling rates are amplified. In contrast, the frequencies where the two sampling rates vote differently will be attenuated as shown in Fig. 7(d).

After unfolding the aliased spectra and recovering the filtered spectrum through voting, we calculate the average signal power in each spike by summing up the voting results and divide it by the spike width. Additionally, we also estimate the average signal power in the spikes using the unfolded spectrum at each sampling rate separately. We classify a band as occupied if all three power estimations in the corresponding spike exceed a pre-selected power threshold. This power threshold is selected based on the noise floor, which is measured when all bands are empty. By using all three estimates, we add hard voting on top of the soft voting which adds more robustness to the occupancy detection.

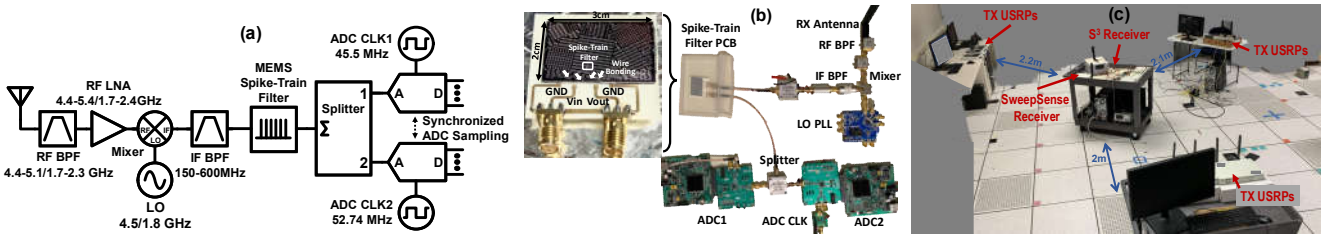

Figure 8: S^3 Prototype System and Evaluation Testbed: (a) S^3 Basic prototype circuit diagram (b) S^3 basic prototype circuit (c) Evaluation testbed.

7 Implementation

A. Basic Prototype: We have built a basic prototype of an S^3 spectrum sensor by combining our MEMS spike-train filter with commodity, off-the-shelf, low-power components. Figure 8(a) shows the circuit diagram of this basic prototype, and the actual prototype is demonstrated in Fig. 8(b). The signal is received through a broadband receiver. It is bandpass filtered and amplified before down-conversion to an intermediate frequency (IF) between 150 and 600 MHz. The IF signal is bandpass filtered and passed through the spike-train filter. It is then split and sampled by the two synchronized ADCs.

We wire-bond the MEMS spike-train filter onto a gold-plated PCB (printed circuit board) as shown in Fig. 8(b). We use K&S 4523A Wedge Bonder and 25 μm Aluminum wire. We use two Anolog Devices LTC2261-14 14-bit ADCs to sample the output of the spike-train filter. This ADC features an 800 MHz wideband input analog bandwidth and low power consumption of 89 mW. The ADC sampling is timed through an external square-wave clock signal. We use the DC1370A ADC evaluation board and the DC890 data acquisition controller to control the ADC sampling through the open-sourced LinearLab Tools Python API. We bypass the input low pass anti-aliasing filter on the ADC evaluation board to maintain the wide analog bandwidth.

B. Extending the Prototype: The above basic prototype using only one spike-train filter can be extended to sense spectra with different center frequency, bandwidth, and channel allocation. Moreover, system level parallelism introduces another degree of freedom and allow us to break the fixed design trade-offs at the filter level.

- **Different Spectrum:** By changing the LO frequency as well as the RF bandpass filter and LNA, we can sense different frequency ranges. In our evaluation, we test at center frequencies of 2.1, 2.4, 4.9, and 5.7 GHz.
- **Larger Bandwidth:** The current spike-train filter supports a bandwidth of 418 MHz. We can extend S^3 to larger bandwidth by either using two sensors and configuring them to sense adjacent spectra or by using two MEMS filters in parallel channels before combining the signals and sampling it as we describe in more detail in appendix A.
- **Narrower Bands:** The spikes in the spike-train filter are separated by 22 MHz. Hence, narrowband signals (< 20 MHz) that are not aligned with the spikes might be filtered out. To address this, we can combine frequency domain sampling with LO frequency sweeping over 22 MHz to capture and sense all the frequencies in the spectrum as shown in our results in section 8. Alternatively, we can design a MEMS filter with narrower spacing as explained in section 3 or use two MEMS filters and set the center frequency to be slightly different as we describe in more detail in appendix A.

C. Testbed: We evaluate the performance on S^3 both through controlled experiments in an indoor wireless testbed as well as through measurements of ambient transmissions outdoors and indoors. The wireless testbed allows us to control the spectrum sparsity, how fast the occupancy changes for different bands, the type of signals transmitted, and the power of various transmissions. It also allows us to know the groundtruth band occupancy in order to evaluate the performance of S^3.

The testbed, shown in Fig. 8(c), can create a 418 MHz spectrum with various occupancy status at different frequencies. It consists of 19 N210 USRP software-defined radios, each transmitting on a 25 MHz bandwidth. While the USRPs are not very far from each other, we vary their transmission power randomly by up to 10 dB and observe received signal SNR that varies by up to 20 dB between different USRP transmitters. To avoid interference from ambient 2.4 and 5 GHz ISM band signals, we conducted experiments in two 418 MHz-wide spectra: 4.73 to 5.15 GHz and 1.93 to 2.35 GHz, each divided into nineteen 22 MHz bands. We vary the spectrum occupancy from 10% to 90%. We also vary the type of modulation being used. We test with single carrier BPSK and QAM as well as OFDM signals. Note that single carrier modulation has a non-flat power spectral density and significantly more leakage, so it results in higher false positive rates as we show in section 8. We also leverage the testbed to compare S^3 with state-of-the-art sensing systems as our baselines. We ran over 5000 experiments with different configuration of occupancy, power, modulation, etc.

8 Results

In this section, we present our main evaluation results along with a few microbenchmarks that provide insights into the performance of S^3 in various spectra.

Evaluation Metrics: We evaluate S^3 using following metrics:

- *False Positive Rate (FPR)*: Percentage of empty bands that S^3 incorrectly reports as occupied.

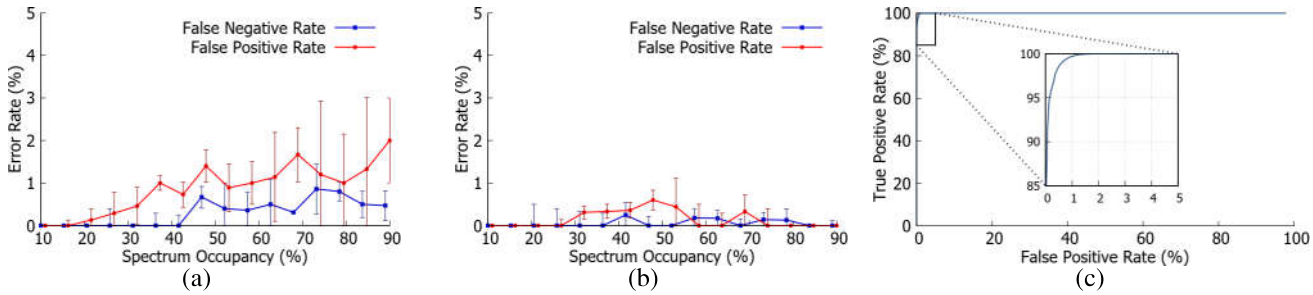

Figure 9: False positives and negatives as a function of spectrum occupancy: The figure shows the false positive rate (FPR) and false negative rate (FNR) of S^3 as the spectrum occupancy increases when: (a) Modulation schemes are randomly picked by transmitters. (b) Transmitters only use OFDM modulation. (c) shows the receiver operating characteristic (ROC) curve of S^3 when the modulation schemes are randomly picked.

- *False Negative Rate (FNR)*: Percentage of occupied bands that S^3 incorrectly reports as empty.
- *True Positive Rate (TNR)*: Percentage of occupied bands that S^3 correctly reports as occupied.

A. Sensing Densely Occupied Spectrum:

Fig. 9 shows S^3's error rate in detecting occupied bands as we vary the total occupancy of the spectrum between 10% and 90%. Fig. 9(a) shows the results when the transmitters randomly pick a modulation scheme (e.g. single carrier BPSK, QAM, or OFDM). In this case, when the total occupancy of the spectrum is less than 30%, S^3 achieves a median false positive rate (FPR) less than 0.5% and a median false negative rate (FNR) of 0%. As the total occupancy increases and the spectrum becomes more crowded, the FPR and FNR gradually increase. However, even when the spectrum is extremely crowded ($\sim$ 90% occupied), S^3 can still achieve 2% median FPR and 0.47% median FNR.

Fig. 9(b) shows the same results when the transmitters only use OFDM modulation. In this case, the FPR and FNR become even smaller at all levels of occupancy with a maximum median FPR of 0.6% and a maximum median FNR of 0.25%. This result can be attributed to two factors: (1) OFDM signals have flat power spectral densities. Therefore, signal power detected in the spike train can more accurately reflect the signal presence in the corresponding channels. (2) Single carrier modulation schemes have lower spectral efficiency and leak power outside their bands, which leads to a higher FPR as can be seen from Fig. 9(a). Finally, Fig. 9(c) shows the receiver operating characteristic (ROC) curve, which demonstrates the trade-off between false positives and false negatives as we vary the threshold for detecting occupied band.

B. Comparison with State-of-the-Art:

We compare S^3 with three baselines from prior work:

- **BigBand:** [21] leverages sparse Fourier transform and uses co-prime sampling to acquire sparse spectrum. It achieves 6$\times$ sub-sampling below the Nyquist rate, but only works when the spectrum is sparse.
- **D-BigBand:** [43] extends BigBand to sense dense spectrum by considering the differential changes in occupancy. It also achieves 6$\times$ sub-sampling, but assumes the changes in the spectrum occupancy over time are sparse.
- **SweepSense:** [20] enhances USRP software-defined radio's ability to quickly scan and sense wideband spectrum. It is able to scan 5 GHz bandwidth in 5 ms with 2×25 MS/s ADC sampling rate.

Table 1: Sum of false positives and negatives for S^3 and State-of-the-Art prior work: The table compares the sum of FPR and FNR of S^3, BigBand, D-BigBand, and SweepSense at different spectrum occupancies.

	BigBand	D-BigBand	SweepSense	S^3
10%	0.38% + 14% (unresolved)	$\sim$ 0.95%	4.88%	0.00%
50%	N/A	$\sim$ 1.75%	13.09%	1.29%
90%	N/A	$\sim$ 3%	13.76%	2.47%

Table 1 shows the sum of FPR and FNR when the total spectrum occupancy is 10%, 50%, and 90%. We compare S^3 directly to the results reported in [21] and [43], because they used custom hardware but were evaluated using the same metrics as ours. One can see that in sparse spectrum (<10% occupied) where BigBand works, BigBand has a total error rate of 0.38% but still cannot recover the status of 14% of the spectrum. In contrast, S^3 accomplishes a 0% error rate at such low spectrum occupancy and samples 8.5$\times$ below the Nyquist rate, which exhibits a 1.4$\times$ gain over BigBand.

D-BigBand is able to work in densely occupied spectrum. It has a total error rate of 0.95% and 3% when the spectrum is 50% and 90% occupied respectively. However, S^3 is able to outperform D-BigBand at all occupancy levels with a 1.2$\times$ to 1.35$\times$ gain in accuracy. Moreover, S^3 also achieves 1.4$\times$ gain in sampling rate reduction and makes no assumptions on the changes in spectrum occupancy. Therefore, unlike D-BigBand, S^3 can monitor highly dynamic spectrum, which we will demonstrate later in this section.

To compare S^3 to SweepSense, we reproduce SweepSense on a N210 USRP with a CBX daughterboard using the codes and FPGA images released by the authors. We use SweepSense to sense the spectrum generated by our testbed along with S^3. In Fig. 10, we show SweepSense's error rate in detecting occupied bands as we vary the spectrum occupancy between 10% and 90%, and when the modulation scheme is randomly picked by the transmitters. This result shows that SweepSense can work in densely occupied spectrum. When

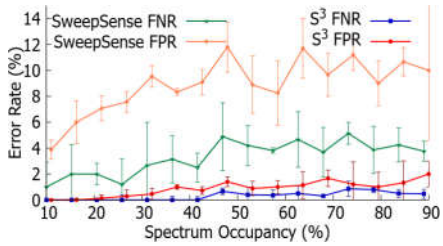

Figure 10: False positives and negatives comparison between S^3 and SweepSense.

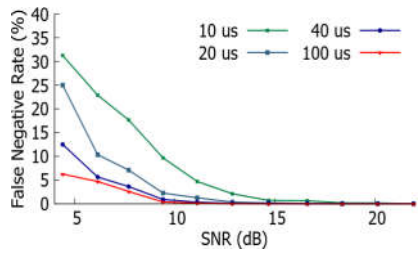

Figure 11: FNR vs SNR and sampling time.

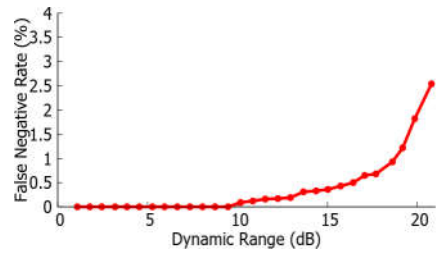

Figure 12: FNR vs dynamic range.

the spectrum is 10% occupied, SweepSense achieves an FPR of 3.88% and an FNR of 1%. As the spectrum becomes more crowded, the FPR and FNR of SweepSense increase, but they remain below 10% and 3.76% respectively even if the spectrum is 90% occupied. SweepSense's higher error rates are likely due to the fact that fast LO sweeping can smear non-empty frequency components, resulting in more leakage from the occupied bands to the adjacent bands, which increases its false positive rates. We also note that, SweepSense requires accurate phase information for the digital chirp demodulation, so it is sensitive to the IQ imbalance in the hardware, which is likely why it underperforms S^3. However, SweepSense is highly valuable as we can combine it with S^3 to capture the power spectral density as we show later in this section.

Next, we present some microbenchmarks that provide more insights into the working of S^3 and its performance.

C. Microbenchmark - Sensitivity:

To understand the ability of S^3 to detect low signal-to-noise ratio (SNR) signals, we examine the FNR of bands with different SNRs. The SNR we show is the average signal power per Hz of RX signal / noise floor. We compare four different sampling duration: 10, 20, 40, and 100 μs. The FNR is high when the SNR is low, however, this can be addressed by increasing the sampling duration. In fact, we can reduce the FNR by 5$\times$ at 3dB SNR. As the SNR gets higher, the FNR goes down and down, and eventually even for shoft sampling windows, the FNR is very low ($\approx 0\%$). Note that $40 \sim 100\ \mu s$ is a short enough window to detect short transient packets and fleeting signals, as it is comparable to the DIFS duration for Wi-Fi carrier sensing (e.g. 34 or 50 μs).

D. Microbenchmark - Dynamic Range:

Here we evaluate the dynamic range of S^3, which is the ratio between the strongest and weakest signal powers S^3 can accurately detect at the same time. It reflects the ability of S^3 to detect low-power signals with the presence of much higher power signals, that would cause interference and lower the signal-to-noise-pluse-interference ratio (SINR), making low-power signals harder to be detected. In Figure 12, we compare the FNRs in experiments with different dynamic ranges. One can see that S^3 achieves very low FNR ($< 0.63\%$) when the power difference between the occupied bands is up to 15 dB. As the signal power difference becomes even larger, the FNR of S^3 increases. When the spectrum dynamic range reaches ~ 21 dB, the FNR of S^3 is 2.54%. This result shows that S^3 can accurately detect (FPN<1%) the relatively weak signals under interference from signals 19 dB stronger.[5]

E. Microbenchmark - Resolving Collisions with Voting:

We want to verify that through voting using two different sampling rates, S^3 can effectively resolve frequency collisions. To this end, we compare S^3 to baselines where we detect the spectrum occupancy using only one ADC. In Fig. 13(a), we qualitatively compare the correctness of occupancy detection on each band in 20 randomly selected experiments. It shows that when using either ADC alone, we have many false positives due to frequency collisions. However, the two ADCs exhibit false positives in different bands, because they experience frequency collisions between different spikes. As a result, through voting S^3 is able to distinguish and resolve false positives where the two sampling rates disagree with each other. Furthermore, we also quantitatively show the FPR of S^3 and baselines. As can be seen from Fig. 13(b), the FPR of S^3 is much lower than those of baselines, which suggests that our voting scheme can effectively leverage the different sampling rates to resolve frequency collisions. Note that ADC1 outperforms ADC2. This is expected because, as we discussed in section 6, 45.5 MHz is an optimized sampling rate that can spread out the spikes and minimize the frequency collisions. In contrast, the second sampling rate of 52.74 MHz is optimized to avoid having the same collisions as ADC1, so it does not work as well by itself.

F. Monitoring Dynamic Spectrum:

S^3 senses all bands in the spectrum in real-time and makes no assumptions on the changes of spectrum occupancy, so it can monitor highly dynamic spectrum with rapidly-changing occupancied bands. To evaluate this ability of S^3 we create a rapidly-changing spectrum in our testbed whose occupied bands change every 327 μs, and as a result, the total spectrum occupancy varies between 0% and 63%. We use S^3 to continuously monitor the occupancy changes in the spectrum, and output a signal power estimation and occupancy detection for every 76 μs-long frame. We show a spectrogram captured by S^3 consisting the signal power detected in every band per

[5] After wire-bonding, the spike-train filter experiences degradation in the out-of-band suppression due to the direct leakage from the input port to the output port of the PCB. Hence, the sensitivity and dynamic range of S^3 also degrade, but this issue can be resolved by better isolation in the PCB design.

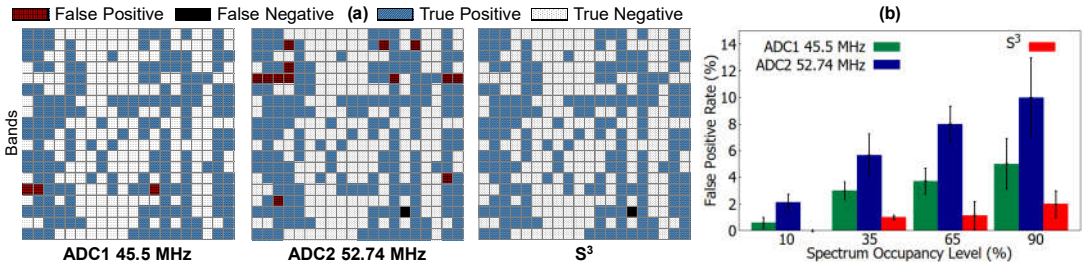

Figure 13: False positives comparison between S^3 and baselines that use single sampling rate: (a) shows the spectrum occupancy detected by S^3 and the baselines in 20 randomly selected experiments. (b) shows the false positive rate comparison as the spectrum occupancy increases.

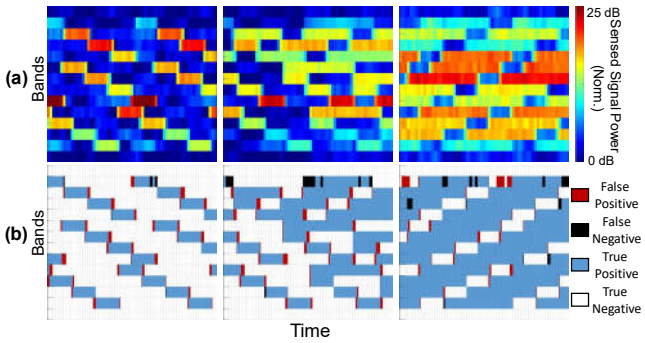

Figure 14: Monitoring rapidly-changing spectrum: The figure shows (a) spectrogram (b) spectrum occupancy captured by S^3 in real time.

frame in Fig. 14(a). Furthermore, we show the accuracy of the corresponding occupancy detection per frame in Fig. 14(b). It shows that S^3 is able to capture the occupancy of rapidly-changing spectrum with great accuracy and time precision.

G. Capturing Wideband Power Spectral Density:

As we mentioned in section 7, we can sweep the LO frequency of S^3 over the 22 MHz spacing between spikes to sense all the frequencies in the spectrum. This enables S^3 to capture the power spectral density (PSD) of the entire wideband spectrum. At every LO frequency, S^3 captures signal power in the spike train and identifies the occupancy of each spike. For the occupied spikes, S^3 uses the signal power estimates in them to reconstruct the PSD at the corresponding RF frequencies. When LO sweeping finishes, all frequencies on the wideband spectrum will be reconstructed. Comparing to conventional spectrum scanners, this extended S^3 prototype only needs to sweep a much narrower frequency range. Therefore, the scanning time is much shorter.

Figure 15(a-c) shows the PSD captured by the extended S^3 prototype of spectra generated by our testbed, along with the detected spike occupancy. We use S^3 and an HP 8563E Spectrum Analyzer to monitor the 1.8 to 2.4 GHz spectrum simultaneously. As one can see, the PSDs captured by S^3 match the ground truth from the spectrum analyzer very well. Besides, we also measure PSDs of real-world spectra, both outdoors and indoors, which are shown in Figure 15(d-f). Figure 15(d) shows the spectrum between 1.8 and 2.4 GHz measured outdoor at our geographical location. It shows that S^3 is able to capture the PSD of 4G LTE signals in Band 2 and 66. In Fig. 15(e) and (f), we show the PSD of 2.4 GHz and 5 GHz Wi-Fi signals captured by S^3 respectively. One can see that Channel 1 and 11 in the 2.4 GHz band as well as four non-overlapping 20 MHz channels (Channel 116, 120, 124, and 128) from 5.57 to 5.65 GHz in the 5 GHz band are being used. Figure 15(d-f) demonstrate that the real-world PSDs captured by S^3 also closely match the spectrum analyzer ground truth. On some frequencies that S^3 classifies as empty, the spectrum analyzer shows some non-zero spectral components. However, this is expected because in our experiments, the spectrum analyzer takes the maximum over a lot more scans than S^3.

9 Related Work

In this section, we provide more related work. For further background, we refer the reader to section 2.

Spectrum sensing has been extensively studied in the past two decades [3, 5, 49]. However, most of this work focuses on narrowband sensing [6, 7, 9, 23, 25, 35, 38, 62]. In contrast, this paper focuses on wideband spectrum sensing to enable dynamic spectrum sharing of many channels. Several systems attempt to sense wideband spectrum using narrowband sensors without sequentially scanning each band [42, 46, 61]. QuickSense [61] leverages analog filters and energy detectors to hierarchically sense wide bandwidth by detecting the total signal power in groups of consecutive channels. However, when the spectrum is densely occupied, QuickSense's approach reduces to sequentially scanning the spectrum. SpecInsight [42] leverages machine learning to predict spectrum occupancy based on learned utilization patterns and optimize which channels to sense. Similarly, [46] uses time-series analysis to predict which bands are occupied. However, these systems are sensitive to training data and assume that transmissions follow predictable patterns. Spectrum sharing is based on opportunistic access and as a result is highly dynamic and unpredictable [50].

Recent work aims to enhance USRP software-defined radio's ability to sense wideband spectrum [20, 26]. SweepSense [20] enables sensing wideband spectrum by quickly sweeping the center frequency of the USRP. It is able to sweep 5 GHz bandwidth in 5 ms, which offers great potential for sensing an extremely wideband spectrum on commercial software radios. However, SweepSense requires

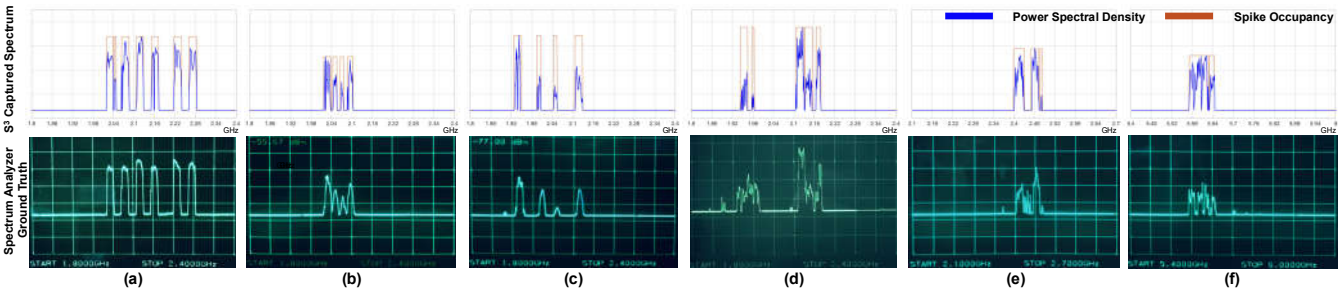

Figure 15: Wideband Power Spectral Density Capture: The figure shows the wideband power spectral density captured by the extended S^3 prototype. Spectrum (a-c) are generated by our evaluation testbed. Spectrum (d-f) are real-world spectra captured both outdoor and indoor.

accurate phase information for the digital chirp demodulation and is sensitive to the IQ imbalance. As a result, our comparison with SweepSense in section 8 shows that it can suffer from a high error rate especially when the spectrum is not sparse. SparSDR [26] reduces the backhaul and computation requirements for sensing sparse spectrum on USRPs, which offers great utility for continuously monitoring underutilized spectra but cannot scale to densely occupied spectrum.

The use of single passband MEMS filters in spectrum sensing has been studied [33, 36]. However, these techniques require an array of channel-select MEMS filters to form a reconfigurable filter bank. In contrast, S^3 only uses a single MEMS spike-train filter that consists of overtone resonators.

Our work is also related to theoretical work on co-prime sampling [54, 56, 57] and multicoset sampling [22, 55] of sparse wideband spectrum. These approaches also do not work for densely occupied spectrum. Moreover, [56, 57] require using k ADCs where k is the number of occupied frequencies. [22, 55] aim to recover the signals in each occupied band and must assume prior knowledge of which bands are occupied. In contrast, S^3 aims to recover the occupancy of each band and uses 2 ADCs irrespective of the number of occupied frequencies. S^3 is further implemented and shown to work in practice.

Sub-Nyquist sampling has been used for test equipment to reconstruct wideband periodic signals [44, 52]. However, these techniques require the signal to be periodic and repeat for a long time in order to take on samples during each period until all samples are recovered. Hence, these techniques are not applicable to real communication signals where the signal is constantly changing and carries different modulated bits.

Finally, some works aim to capture spectrum usage at large geographical and time scales through crowdsourcing [10, 37, 63]. S^3 is complementary to these works, as it enables real-time wideband occupancy detection of every single sensor with minimum data size and computational complexity.

10 Limitations

In this section, we discuss some limitations of S^3.

- The frequency-domain sampling rate and maximum sensing bandwidth is limited by the filter design trade-offs. As a result, narrowband signals (< 20 MHz) and over GHz-wide spectrum cannot be sensed using a single spike-train filter. This can be resolved by hopping the LO frequency as shown in Section 8. Alternatively, we can also use the extended architectures proposed in appendix A that combines multiple spike-train filters in parallel to break the fixed filter-level design trade-offs.
- Sub-Nyquist sampling leads to aliasing of both signals and noise, which typically lowers the signal SNR and degrades the spectrum sensor's sensitivity. To minimize the loss of SINR, we design the spike-train filters to have low insertion loss and high out-of-band suppression i.e., most of the noise is filtered out before it aliases. Moreover, instead of detecting signal power, known signals like the preambles can be leveraged to improve the sensitivity [48]. However, directly applying this technique to S^3 would require further research as the preambles might become corrupted after applying the filter.
- While S^3 can detect the occupancy of the different bands and reconstruct the power spectral density of the spectrum, it cannot recover complex I and Q samples of the signal. As results, S^3 cannot reconstruct the signal itself or decode the data in the signal.

11 Conclusion

This paper presents S^3, a new efficient real-time wideband spectrum sensing mechanism that can work in densely occupied spectrum. S^3 monitors only a small fraction of bandwidth in each band to accomplish significantly below-Nyquist sampling and hence great energy efficiency. It leverages recent advances in RF MEMS filtering solution that enables sampling the spectrum along the frequency axis. Empirical evaluation demonstrates that S^3 can accurately sense densely occupied spectrum and rapidly-changing spectrum; we also show that S^3 can be extended to capture the power spectral density of the entire spectrum. We believe S^3 can enable dynamic spectrum access and very high spectrum utilization.

Acknowledgments: We would like to thank our shepherd Dinesh Bharadia as well as the reviewers for their feedback. This work is funded in part by NSF Award 1824320.

References

[1] O. Abari, F. Chen, F. Lim, and V. Stojanović. Performance trade-offs and design limitations of analog-to-information converter front-ends. In *2012 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 5309–5312. IEEE, 2012.

[2] O. Abari, F. Lim, F. Chen, and V. Stojanović. Why analog-to-information converters suffer in high-bandwidth sparse signal applications. *IEEE Transactions on Circuits and Systems I: Regular Papers*, 60(9):2273–2284, 2013.

[3] A. Ali and W. Hamouda. Advances on spectrum sensing for cognitive radio networks: Theory and applications. *IEEE Communications Surveys Tutorials*, 19(2):1277–1304, 2017.

[4] Analog Devices, Norwood, MA, USA. *12-Bit, 2.6 GSPS/2.5 GSPS/2.0 GSPS, 1.3 V/2.5 V Analog-to-Digital Converter AD9625 Data Sheet*.

[5] Y. Arjoune and N. Kaabouch. A comprehensive survey on spectrum sensing in cognitive radio networks: Recent advances, new challenges, and future research directions. *Sensors (Basel, Switzerland)*, 19(1):126, 01 2019.

[6] E. Axell and E. G. Larsson. Optimal and sub-optimal spectrum sensing of ofdm signals in known and unknown noise variance. *IEEE Journal on Selected Areas in Communications*, 29(2):290–304, Feb. 2011.

[7] P. Bahl, R. Chandra, T. Moscibroda, R. Murty, and M. Welsh. White space networking with wi-fi like connectivity. *ACM SIGCOMM Computer Communication Review*, 39(4):27–38, 2009.

[8] M. Bao and H. Yang. Squeeze film air damping in mems. *Sensors and Actuators A: Physical*, 136(1):3 – 27, 2007. 25th Anniversary of Sensors and Actuators A: Physical.

[9] D. Bhargavi and C. R. Murthy. Performance comparison of energy, matched-filter and cyclostationarity-based spectrum sensing. In *2010 IEEE 11th International Workshop on Signal Processing Advances in Wireless Communications (SPAWC)*, pages 1–5, June 2010.

[10] A. Chakraborty, M. S. Rahman, H. Gupta, and S. R. Das. Specsense: Crowdsensing for efficient querying of spectrum occupancy. In *IEEE INFOCOM 2017-IEEE Conference on Computer Communications*, pages 1–9. IEEE, 2017.

[11] A. A. Cheema and S. Salous. Digital fmcw for ultrawideband spectrum sensing. *Radio Science*, 51(8):1413–1420, 2016.

[12] DigiKey. Data acquisition - analog to digital converters (adc). https://www.digikey.com/products/en/integrated-circuits-ics/data-acquisition-analog-to-digital-converters-adc, Sep. 2020.

[13] Federal Communications Commission. 3.5 ghz band overview. https://www.fcc.gov/wireless/bureau-divisions/mobility-division/35-ghz-band/35-ghz-band-overview, Jan. 2020.

[14] Federal Communications Commission. Fcc adopts new rules for the 6 ghz band, unleashing 1,200 megahertz of spectrum for unlicensed use. https://www.fcc.gov/document/fcc-opens-6-ghz-band-wi-fi-and-other-unlicensed-uses, Apr. 2020.

[15] Federal Communications Commission. White space. https://www.fcc.gov/general/white-space, Jan. 2020.

[16] Y. Ghasempour, C. R. C. M. da Silva, C. Cordeiro, and E. W. Knightly. Ieee 802.11ay: Next-generation 60 ghz communication for 100 gb/s wi-fi. *IEEE Communications Magazine*, 55(12):186–192, 2017.

[17] B. Ghazi, H. Hassanieh, P. Indyk, D. Katabi, E. Price, and L. Shi. Sample-optimal average-case sparse fourier transform in two dimensions. In *2013 51st Annual Allerton Conference on Communication, Control, and Computing (Allerton)*, pages 1258–1265. IEEE, 2013.

[18] S. Gong, Y. Song, T. Manzaneque, R. Lu, Y. Yang, and A. Kourani. Lithium niobate mems devices and subsystems for radio frequency signal processing. In *2017 IEEE 60th International Midwest Symposium on Circuits and Systems (MWSCAS)*, pages 45–48, 2017.

[19] Y. M. Greshishchev, J. Aguirre, M. Besson, R. Gibbins, C. Falt, P. Flemke, N. Ben-Hamida, D. Pollex, P. Schvan, and S. Wang. A 40gs/s 6b adc in 65nm cmos. In *2010 IEEE International Solid-State Circuits Conference - (ISSCC)*, pages 390–391, 2010.

[20] Y. Guddeti, R. Subbaraman, M. Khazraee, A. Schulman, and D. Bharadia. Sweepsense: Sensing 5 ghz in 5 milliseconds with low-cost radios. In *16th USENIX Symposium on Networked Systems Design and Implementation (NSDI 19)*, Feb. 2019.

[21] H. Hassanieh, L. Shi, O. Abari, E. Hamed, and D. Katabi. Ghz-wide sensing and decoding using the fourier transform. In *IEEE INFOCOM 2014 - IEEE Conference on Computer Communications*, Apr. 2014.

[22] C. Herley and Ping Wah Wong. Minimum rate sampling and reconstruction of signals with arbitrary frequency support. *IEEE Transactions on Information Theory*, 45(5):1555–1564, 1999.

[23] S. Hong and S. Katti. Dof: a local wireless information plane. In *Proceedings of the ACM SIGCOMM 2011 conference*, pages 230–241, 2011.

[24] M. Kadota, S. Tanaka, Y. Kuratani, and T. Kimura. Ultrawide band ladder filter using sh0 plate wave in thin linbo3 plate and its application. In *2014 IEEE International Ultrasonics Symposium*, pages 2031–2034, 2014.

[25] S. Kapoor, S. Rao, and G. Singh. Opportunistic spectrum sensing by employing matched filter in cognitive radio network. In *2011 International Conference on Communication Systems and Network Technologies*, pages 580–583, June 2011.

[26] M. Khazraee, Y. Guddeti, S. Crow, A. C. Snoeren, K. Levchenko, D. Bharadia, and A. Schulman. Sparsdr: Sparsity-proportional backhaul and compute for sdrs. In *Proceedings of the 17th Annual International Conference on Mobile Systems, Applications, and Services*, MobiSys '19, page 391–403, 2019.

[27] J. N. Laska, W. F. Bradley, T. W. Rondeau, K. E. Nolan, and B. Vigoda. Compressive sensing for dynamic spectrum access networks: Techniques and tradeoffs. In *2011 IEEE International Symposium on Dynamic Spectrum Access Networks (DySPAN)*, 2011.

[28] R. Lu, T. Manzaneque, Y. Yang, A. Kourani, and S. Gong. Lithium niobate lateral overtone resonators for low power frequency-hopping applications. In *2018 IEEE Micro Electro Mechanical Systems (MEMS)*, pages 751–754, 2018.

[29] R. Lu, T. Manzaneque, Y. Yang, J. Zhou, H. Hassanieh, and S. Gong. Rf filters with periodic passbands for sparse fourier transform-based spectrum sensing. *Journal of Microelectromechanical Systems*, 27(5):931–944, 2018.

[30] Y. Ma, Y. Gao, A. Cavallaro, C. G. Parini, W. Zhang, and Y. Liang. Sparsity independent sub-nyquist rate wideband spectrum sensing on real-time tv white space. *IEEE Transactions on Vehicular Technology*, 66(10):8784–8794, 2017.

[31] M. Mishali and Y. C. Eldar. From theory to practice: Sub-nyquist sampling of sparse wideband analog signals. *IEEE Journal of Selected Topics in Signal Processing*, Apr. 2010.

[32] M. Mishali, Y. C. Eldar, O. Dounaevsky, and E. Shoshan. Xampling: Analog to digital at sub-nyquist rates. *IET Circuits, Devices Systems*, Jan. 2011.

[33] T. Mukherjee, G. K. Fedder, H. Akyol, U. Arslan, J. Brotz, F. Chen, A. Jajoo, C. Lo, A. Oz, D. P. Ramachandran, V. K. Saraf, M. Sperling, and J. Stillman. Reconfigurable mems-enabled rf circuits for spectrum sensing. In *Government Microcircuit Applications and Critical Technology Conference*, 2005.

[34] B. Murmann. A/d converter trends: Power dissipation, scaling and digitally assisted architectures. In *2008 IEEE Custom Integrated Circuits Conference*, pages 105–112, 2008.

[35] A. Nasser, A. Mansour, K. C. Yao, H. Charara, and M. Chaitou. Efficient spectrum sensing approaches based on waveform detection. In *The Third International Conference on e-Technologies and Networks for Development (ICeND2014)*, pages 13–17, April 2014.

[36] C. T. C. Nguyen. Mems-based rf channel selection for true software-defined cognitive radio and low-power sensor communications. *IEEE Communications Magazine*, 51(4):110–119, 2013.

[37] D. Pfammatter, D. Giustiniano, and V. Lenders. A software-defined sensor architecture for large-scale wideband spectrum monitoring. In *Proceedings of the 14th International Conference on Information Processing in Sensor Networks*, pages 71–82, 2015.

[38] H. Rahul, N. Kushman, D. Katabi, C. Sodini, and F. Edalat. Learning to share: narrowband-friendly wideband networks. *ACM SIGCOMM Computer Communication Review*, 38(4):147–158, 2008.

[39] M. Rashidi, K. Haghighi, A. Panahi, and M. Viberg. A nlls based sub-nyquist rate spectrum sensing for wideband cognitive radio. In *2011 IEEE International Symposium on Dynamic Spectrum Access Networks (DySPAN)*, pages 545–551. IEEE, 2011.

[40] M. Rinaldi, C. Zuniga, C. Zuo, and G. Piazza. Ultra-thin super high frequency two-port aln contour-mode resonators and filters. In *TRANSDUCERS 2009 - 2009 International Solid-State Sensors, Actuators and Microsystems Conference*, pages 577–580, 2009.

[41] S. K. Sharma, E. Lagunas, S. Chatzinotas, and B. Ottersten. Application of compressive sensing in cognitive radio communications: A survey. *IEEE Communications Surveys Tutorials*, 18(3):1838–1860, 2016.

[42] L. Shi, P. Bahl, and D. Katabi. Beyond sensing: Multi-ghz realtime spectrum analytics. In *12th USENIX Symposium on Networked Systems Design and Implementation (NSDI 15)*, May 2015.

[43] L. Shi, H. Hassanieh, and D. Katabi. 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, page 13–16, 2014.

[44] A. J. Silva. Reconstruction of undersampled periodic signals. *MIT Technical Report*, 1986.

[45] Y. Song and S. Gong. Wideband spurious-free lithium niobate rf-mems filters. *Journal of Microelectromechanical Systems*, 26(4):820–828, 2017.

[46] J. Su and W. Wu. Wireless spectrum prediction model based on time series analysis method. In *Proceedings of the 2009 ACM workshop on Cognitive radio networks*, pages 61–66, 2009.

[47] S. Subramaniam, H. Reyes, and N. Kaabouch. Spectrum occupancy measurement: An autocorrelation based scanning technique using usrp. In *2015 IEEE 16th Annual Wireless and Microwave Technology Conference (WAMICON)*, pages 1–5, 2015.

[48] H. Sudo, K. Kosaka, H. Kanemoto, N. Gejoh, T. Yasunaga, and M. Uesugi. Study of spectrum sensing scheme using received power within preamble signals. In *2017 20th International Symposium on Wireless Personal Multimedia Communications (WPMC)*, pages 592–597, 2017.

[49] H. Sun, A. Nallanathan, C. Wang, and Y. Chen. Wideband spectrum sensing for cognitive radio networks: a survey. *IEEE Wireless Communications*, 20(2):74–81, April 2013.

[50] The Software Defined Radio Forum Inc. Application of management technologies in dynamic spectrum sharing, Jul. 2019.

[51] J. A. Tropp, J. N. Laska, M. F. Duarte, J. K. Romberg, and R. G. Baraniuk. Beyond nyquist: Efficient sampling of sparse bandlimited signals. *IEEE transactions on information theory*, 56(1):520–544, 2009.

[52] N. Tzou, D. Bhatta, S. Hsiao, H. W. Choi, and A. Chatterjee. Low-cost wideband periodic signal reconstruction using incoherent undersampling and back-end cost optimization. In *2012 IEEE International Test Conference*, pages 1–10, 2012.

[53] U.S. Government. CFR title 47 section 96.67 environmental sensing capability, Jan. 2020.

[54] P. P. Vaidyanathan and P. Pal. Sparse sensing with coprime samplers and arrays. *IEEE Transactions on Signal Processing*, 59(2):573–586, 2011.

[55] R. Venkataramani and Y. Bresler. Perfect reconstruction formulas and bounds on aliasing error in sub-nyquist nonuniform sampling of multiband signals. *IEEE Transactions on Information Theory*, 46(6):2173–2183, 2000.

[56] X. G. Xia. On estimation of multiple frequencies in undersampled complex valued waveforms. *IEEE transactions on signal processing*, 47(12):3417–3419, 1999.

[57] X. G. Xia. An efficient frequency-determination algorithm from multiple undersampled waveforms. *IEEE Signal Processing Letters*, 7(2):34–37, 2000.

[58] R. T. Yazicigil, T. Haque, M. R. Whalen, J. Yuan, J. Wright, and P. R. Kinget. Wideband rapid interferer detector exploiting compressed sampling with a quadrature analog-to-information converter. *IEEE Journal of Solid-State Circuits*, Dec 2015.

[59] J. Yoo, S. Becker, M. Loh, M. Monge, E. Candes, and A. Emami-Neyestanak. A 100mhz–2ghz 12.5 x sub-nyquist rate receiver in 90nm cmos. In *2012 IEEE Radio Frequency Integrated Circuits Symposium*, pages 31–34. IEEE, 2012.

[60] J. Yoo, C. Turnes, E. B. Nakamura, C. K. Le, S. Becker, E. A. Sovero, M. B. Wakin, M. C. Grant, J. Romberg, A. Emami-Neyestanak, and E. Candes. A compressed sensing parameter extraction platform for radar pulse signal acquisition. *IEEE Journal on Emerging and Selected Topics in Circuits and Systems*, 2(3):626–638, 2012.

[61] S. Yoon, L. E. Li, S. C. Liew, R. R. Choudhury, I. Rhee, and K. Tan. Quicksense: Fast and energy-efficient channel sensing for dynamic spectrum access networks. In *2013 Proceedings IEEE INFOCOM*, April 2013.

[62] T. Yucek and H. Arslan. A survey of spectrum sensing algorithms for cognitive radio applications. *IEEE communications surveys & tutorials*, 11(1):116–130, 2009.

[63] Y. Zeng, V. Chandrasekaran, S. Banerjee, and D. Giustiniano. A framework for analyzing spectrum characteristics in large spatio-temporal scales. In *The 25th Annual International Conference on Mobile Computing and Networking*, pages 1–16, 2019.

[64] C. Zuo, N. Sinha, and G. Piazza. Very high frequency channel-select mems filters based on self-coupled piezoelectric aln contour-mode resonators. *Sensors and Actuators A: Physical*, 160(1):132 – 140, 2010.

Appendix A Alternative S^3 Architectures

As we mentioned in section 7, we can extend S^3 to use two or more MEMS filters in parallel to enlarge the sensing bandwidth or enable S^3 to sense different channel allocations including narrower bands and even non-uniformly allocated

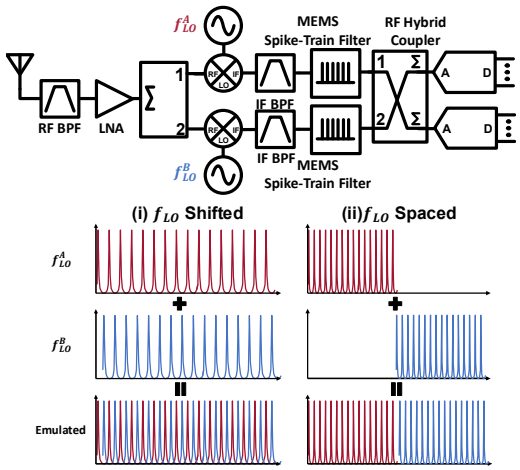

Figure 16: Circuit diagram and emulated spike-train filter for alternative S^3 architecture leveraging LO center frequency difference.

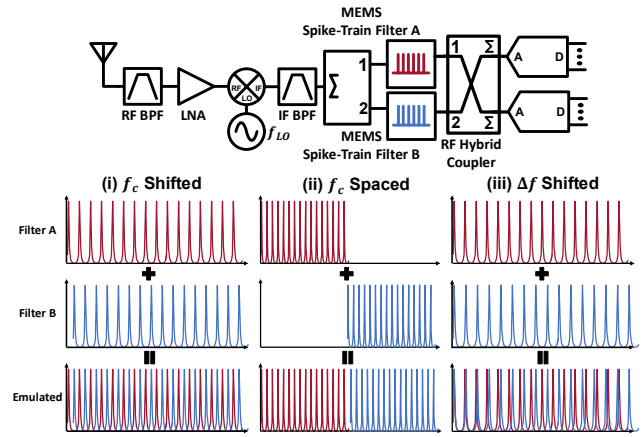

Figure 17: Circuit diagram and emulated spike-train filter for alternative S^3 architecture leveraging different MEMS spike-train filters.

bands. This system level parallelism introduces another degree of freedom and allows us to break the fixed design tradeoffs at the filter level. Here, we present some alternative architectures of S^3 that combine two spike-trains filters.

(1) Changing LO Center Frequencies: First, we can combine two identical spike-train filters that are fabricated using the same process. Hence, these two filters will have almost identical frequency responses with the same center frequency f_c and spacing between spikes Δf. In order to cover different frequencies in the RF spectrum, we use the two filters on separate receiver RF chains with different LO frequencies. We demonstrate the circuit diagram and the emulated spike trains in the RF spectrum in Fig. 16.

After bandpass filtering and amplifying the received signal, we split the RF signal into two channels, and use two LOs with center frequencies f_{LO}^A and f_{LO}^B to down-convert the signal to the IF frequencies. Then we pass each IF signal into a spike-train filter to sampling the spectrum along the frequency axis. Based on the LO frequency difference $df_{LO} = f_{LO}^B - f_{LO}^A$, we can emulate two types of spike trains, as shown in Fig. 16(i) and (ii). When $\Delta f_{LO} < \Delta f$, the two spike trains are slightly shifted on the frequency axis as shown in Fig. 16(i). As a result, we can emulate a spike-train filter with narrower spacing between the spikes. This increases the frequency domain sampling rate of the filter and enables S^3 to sense narrower channel bandwidths. On the other hand, when $df_{LO} = K\Delta f$, the two spike trains are concatenated along the frequency axis as shown in Fig. 16(ii). In this way, a longer spike train with more spikes covering wider bandwidth is emulated.

Although it is straight forward to sample the two IF signals separately, the number of ADCs required will increase linearly with the number of spike-train filters. Instead, after passing IF signals on the two channels through the spike-train filters, we combine the filtered signals and sample the combined signal using two low-speed ADCs. The analog combination and splitting can be achieved using an RF power combiner in series with an RF power splitter, but it can also be done using a RF hybrid coupler. Note that with more spikes in the emulated filter, there will be more aliasing in the subsampled spectrum. Therefore, a higher ADC sampling rate might be needed, but the sampling rate should be able to scale sublinearly with respect to the number of spikes. Besides, the ADC input cutoff frequency needs to be higher than the spike-train bandwidth.

The advantage of this architecture is that we can use the same MEMS spike-train filter on the two channels without needing to redesign a new filter. However, it requires two LOs and mixers which increases the cost and power consumption of the system. [6]

(2) Changing Spike-Train Filter Structure: Instead of introducing a second local oscillator, we can use only one LO and two different spike-train filters to emulate spike trains with wider bandwidth as well as narrower or nonuniform spike spacing. As we mentioned in section 3, we can modify the width of the piezoelectric film and the position of the electrodes to obtain different Δf and f_c. When two spike-train filters with different Δf and/or f_c are combined in parallel, we can emulate a spike train with more sophisticated sparsity structures. We show the circuit diagram for this type of alternative architectures in Fig. 17, along with three emulated spike-train filter frequency responses.

In this architecture, the down-converted IF signal is split and filtered by two different MEMS spike-train filters, whose center frequencies and spike spacing are f_c^A, f_c^B, and Δf^A, Δf^B. The output spectra of the two filters are then combined and sampled by two low-speed ADCs. Using this architecture, we can emulate the same spike trains as the first alternative architecture. For instance, when the difference between the filter center frequencies $df_c = f_c^B - f_c^A < \Delta f$, as shown in Fig. 17(i), the two spike trains are slightly shifted on the

[6] Since the power consumption of spectrum sensors is dictated by the ADC [58], the additional power consumption of the second LO is not the primary concern.

frequency axis and emulate a spike-train filter with narrower spacing between the spikes. Besides, when $df_c = K\Delta f$, the two spike trains are spaced by the bandwidth of the spike train and emulate a wider bandwidth spike train as shown in Fig. 17(ii). However, in additional to enlarging the filter bandwidth and narrowing the spike spacing, we can even emulate a non-uniformly spike train as as shown in Fig. 17(iii). This is achieved by combining two spike-train filters with different Δf^A and Δf^B. Such spike train profile provides us with all sorts of frequency resolutions across the spectrum to accommodate the different channel bandwidth required by the secondary users in TV Whitespace and CBRS bands.