Yeswanth Guddeti, Raghav Subbaraman, Moein Khazraee, Aaron Schulman and Dinesh Bharadia *UC San Diego*

Editors: Nic Lane and Xia Zhou

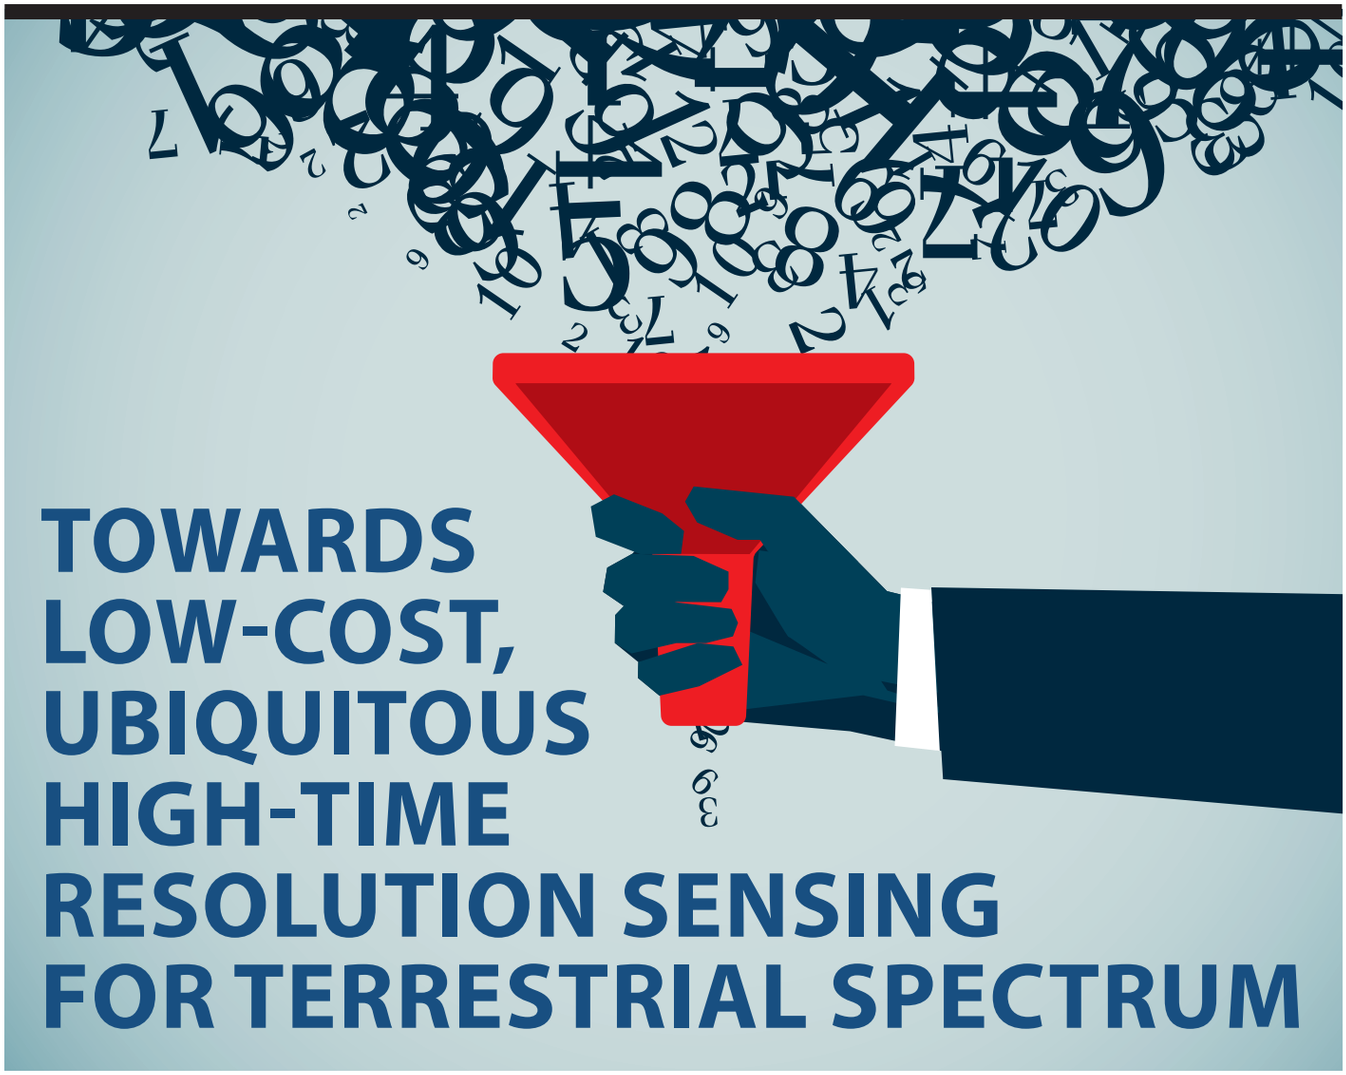

TOWARDS LOW-COST, UBIQUITOUS HIGH-TIME RESOLUTION SENSING FOR TERRESTRIAL SPECTRUM

Excerpted from "SweepSense: Sensing 5 GHz in 5 Milliseconds with Low-cost Radios," from *Proceedings of the 16th USENIX Symposium on Networked Systems Design and Implementation (NSDI '19)*, with permission. https://www.usenix.org/conference/nsdi19/presentation/guddeti © USENIX 2019

Spectrum utilization has become increasingly fragmented and diverse, with millions of devices supporting multiple protocols and use cases. A smartphone today uses LTE, Bluetooth, and Wi-Fi simultaneously, spanning multiple GHz of spectrum. In such an environment, comprehensive knowledge of spectrum usage can help in ushering in better networked system design, unified control, and provide new ways of understanding cyber-physical interactions. A versatile spectrum sensor that can monitor all utilized bands quickly and at low cost is a crucial enabler for such applications. Such a sensor is difficult to realize due to the vast size of the terrestrial spectrum and the necessity to keep up with fleeting wireless communication signals in innumerable bands. Additionally, it needs to be inexpensive and suitable for mass deployment. We introduce a new paradigm, in which the spectrum sensor can "sweep" over a wide band of spectrum quickly and extract useful information, such as occupancy and protocol type. The system, dubbed *SweepSense*, is prototyped on inexpensive hardware and opens up opportunities for new applications that were previously intractable, even with expensive solutions.

Illustration, istockphoto.com

Recent applications in wireless and networked system deployment have created an unprecedented requirement for practical spectrum sensing. For example, the FCC granted permission for LTE providers to share licensed spectrum in the 3.5 GHz Citizens Broadband Radio Service (CBRS) band with military radars, provided that channel access is managed using sensors. These sensors are required to detect millisecond-long military radar bursts anywhere within the 150 MHz of assigned bandwidth [1]. The adoption of CBRS is a prelude to larger-scale dynamic spectrum management and coexistence schemes with more stringent requirements on sensing.

To cater to these applications, we need a platform that has high time resolution and can detect signals whose lifespans are of the order of milliseconds. Further, for versatile application and rich data association, the platform should be able to sustain this resolution across a wide-bandwidth of interest in the terrestrial spectrum spanning multiple GHz. Unfortunately, only complex and expensive spectrum sensors satisfy both the established criteria. For example, some radios can sample several GHz of RF bandwidth continuously (e.g., OneRadio [2]). However, they are expensive (~$175,000) due to their high-speed Analog-to-Digital Converters (ADC); and complex due to the massive computational power needed to perform real-time signal processing on high sample rates (e.g., GPUs or FPGAs). Test equipment like RF spectrum analyzers can also scan GHz wide spectrum [3], but are designed to measure the absolute power of a transmitter accurately (e.g., for certification), or discover bugging devices operating in esoteric bands. Lack of time-domain samples limit the signal analysis capability of spectrum analyzers, and their cost prohibits widespread adoption.

Conventional narrow-bandwidth (<50 MHz) radios, such as Software Defined Radios (SDR) (e.g., USRP or or HackRF) are relatively inexpensive, but cannot observe GHz of frequency at once. The sensing bandwidth of these radios can be improved by intelligent tuning. [6] However, they are still likely to miss transmission due to their narrow bandwidth and the downtime that they experience during tuning (as shown in Figure 1). Some authors have also proposed sub-Nyquist approaches like Sparse FFT,

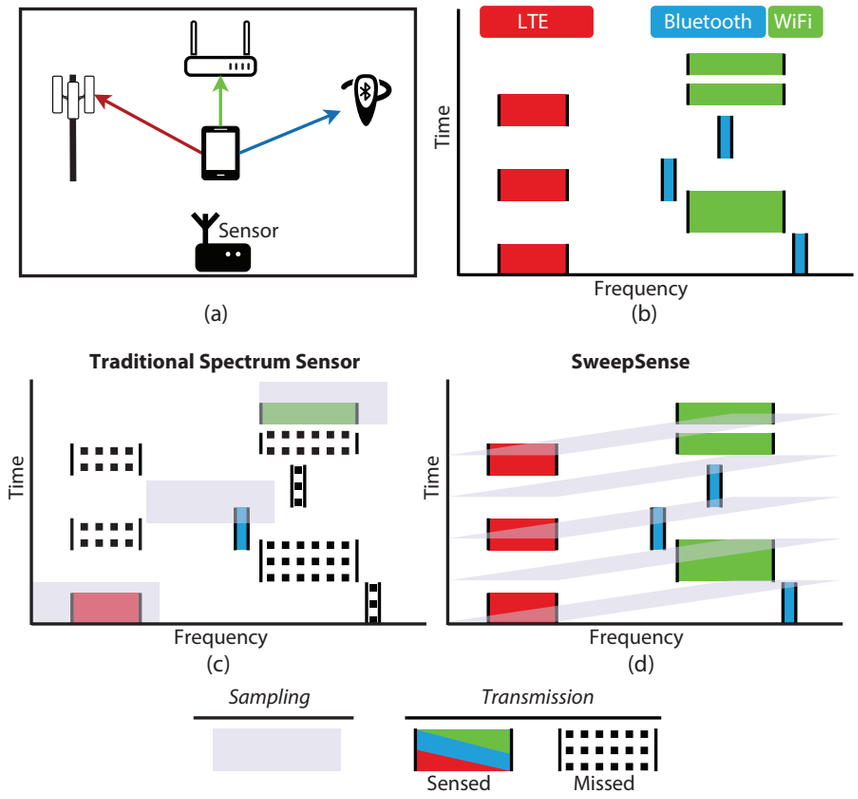

FIGURE 1. (a) A common scenario where a single device may communicate on frequency bands with different protocols. Here we consider LTE, Bluetooth and Wi-Fi with a spectrum sensor in the environment; (b) A waterfall plot of the ground truth spectrum usage by each signal that the device emits, color coded to highlight their characteristics in frequency and time; (c) A traditional spectrum sensor with limited bandwidth has to tune to a band, capture some samples, and move to the next one with some re-tuning delay. Note that many transmissions are missed; (d) SweepSense rapidly sweeps its center frequency, rather than iteratively tuning and capturing the transmissions one frequency at a time, allowing it to sense a part of every transmission.

compressed-sensing [7] to increase sensing bandwidth while using a low bandwidth ADC to sample the entire spectrum. Aside from the assumption on spectrum usage sparsity, these systems directly expose the ADC to power from every band. A single strong signal could therefore saturate the system and affect its usability. Naturally, one can ask the question: *Is there a way to scan full-terrestrial spectrum on inexpensive radios while not missing packets in any band?*

DESIGN OF SweepSense

Our work focuses on answering this question by combining the best aspects of conventional spectrum analyzers and low-cost SDRs. We introduce a new paradigm in spectrum sensing, called SweepSense, which achieves both wide sensing bandwidth and high time resolution with off-the-shelf narrow-bandwidth radios. We observe that sensing a wireless signal shouldn't require us to see it in its entirety; its characteristic features should be identifiable within a limited time-frame. Further, the most natural way to observe signals is to tune to its band and suppress other signals outside the region of interest. Motivated by these two observations, we present a fundamental shift in the receiver architecture for sensing radios. Instead of tuning to each frequency, sampling for a short time, then switching to the next frequency, SweepSense rapidly sweeps the frequency of the receiver across the spectrum (Figure 1). By sweeping quickly across the spectrum and capturing parts of every band, SweepSense achieves high time resolution with a narrow bandwidth radio.

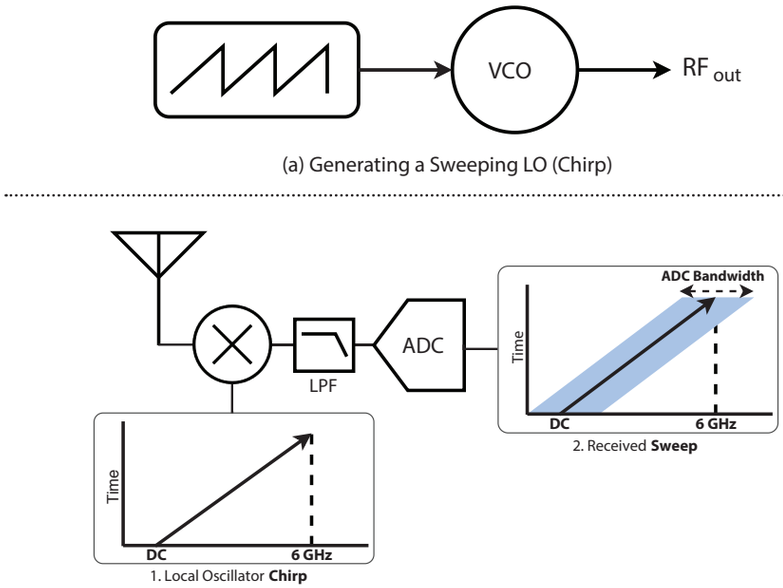

WE INTRODUCE A NEW PARADIGM, IN WHICH THE SPECTRUM SENSOR CAN "SWEEP" OVER A WIDE BAND OF SPECTRUM QUICKLY AND EXTRACT USEFUL INFORMATION, SUCH AS OCCUPANCY AND PROTOCOL TYPE.

FIGURE 2. (a) Generating a sweeping LO frequency using a VCO; (b) Using an LO with increasing frequency to downconvert allows sweeping of spectrum.

Note that baseband I/Q sampling is done while the radio sweeps in frequency. This sets it apart from spectrum analyzers in the sense that it can provide an instantaneous snapshot of actual spectrum, allowing for deeper insights than just power measurement. If the system can sweep at a rate that it can cycle through all bands of interest faster than the time between successive transmissions, then it would have effectively missed no transmissions. Unlike compressed sensing, SweepSense only captures a small amount of spectrum at any point of time avoiding ADC saturation due to strong signals. Such a tool for spectrum monitoring can open up new opportunities in various impact areas, two of which are discussed in the "Evaluation and Results" section (page 26).

There are, however, several challenges that we must overcome to demonstrate that SweepSense is practical and feasible. First, off-the-shelf radios are not designed to sweep. SweepSense is impactful only if it can be implemented on low-cost hardware, enabling distributed deployment. Second, since the center frequency of a sweeping radio changes with time, the samples that it provides are distorted. Therefore, conventional signal processing techniques, developed almost exclusively for fixed frequency receivers, cannot be directly applied to SweepSense samples. Added to these distortions is the fact that a sweeping radio captures only a small part of every transmission, leaving us with the task of making inferences on limited data. In the rest of this article, we describe a prototype SweepSense implementation on an inexpensive SDR platform and validate our ideas with practical evaluations. We also go over our technique to remove sweep-induced distortions and methods to obtain useful information from limited samples using cyclostationary analysis.

MAKING OFF-THE-SHELF RADIOS SWEEP

Conventional radios introduced in wireless systems classes tune to a passband frequency of interest, downconvert them to a lower baseband frequency, and then sample them using an ADC with a low-pass response. The downconversion process typically uses a mixer fed with a Local Oscillator (LO) generated by a Voltage Controlled Oscillator (VCO). The LO frequency determines the specific passband frequency that downconverts to the baseband. Therefore, sweeping the LO frequency can allow the radio to sweep through passband frequencies. We use the VCO in open-loop, as shown in Figure 2(a). Providing a sawtooth voltage waveform to the input of a VCO causes the output to ramp up in frequency. The output of the VCO is directly fed to the mixer; the architecture is shown in Figure 2(b). The Low-Pass Filter (LPF) ensures that the ADC is exposed only to a predefined bandwidth around the instantaneous frequency of the LO, effectively mitigating saturation issues. However, using the VCO in open-loop requires breaking a feedback loop and introduces several new challenges, and is addressed next.

UNSWEEPING: REMOVING DISTORTIONS

Recall that baseband samples obtained while using a sweeping LO to downconvert cannot be treated the same way as fixed-frequency downconversion samples. Additionally, the open-loop chirp generation architecture doesn't guarantee a direct linear relationship with the input voltage and the output frequency. All we can be sure of is that the output frequency will monotonically increase as we ramp the input voltage. To allow the application of standard signal processing techniques on SweepSense samples, we develop an "unsweeping" technique that addresses the challenges above. Intuitively, each sample was obtained at a different center frequency; if we know this time-varying center frequency, we can then account for it in post-processing and remove it. Indeed, we show that this is possible through a two-step process:

Calibration: This is a one-time step to ascertain how the time-varying center frequency behaves while sweeping. A known single-frequency tone is injected into the system, and baseband samples captured while sweeping. These samples are the required calibration samples.

Recovery: After performing a capture using SweepSense, the baseband signal is multiplied sample-by-sample with the inverse of the calibration samples to remove distortions introduced due to sweeping.

This simple two-step process works without making any assumptions on the nature of the ramping frequency.

INFERENCE ON LIMITED DATA

Extending the use of SweepSense beyond just power detection to applications like protocol classification requires processing and subsequent inferencing on the unswept samples. However, we only get a small portion of every transmission due to the sweeping effect. Therefore, many conventional methods, like decoding and preamble detection, cannot be applied. We build insights from the field of cyclostationary analysis [8], developed over the foundation that all human-made signals have repeating patterns. A corollary to the above statement is that every protocol has correlations in time and frequency. These patterns are agnostic to the specific constructions of signals like preambles or actual modulated data and rely only on underlying repeating patterns. By amplifying these signatures through time-domain and spectral correlations, it is possible to differentiate between protocols and even infer some useful knowledge about them, like symbol time and modulation schemes. We show in the extended work ([9]) that unswept SweepSense samples preserve the repeating features of the original signal, allowing the use of the extensive repertoire of cyclostationary signal analysis out of the box.

EVALUATION AND RESULTS

We built and tested SweepSense on a USRP N210 platform using modified CBX and SBX daughtercards. The architecture is generalizable to other software radios where the synthesizer/VCO is accessible for modification. The AUX-DAC on the USRP motherboard is used to generate the sawtooth voltage waveform that drives the VCO. The CBX daughtercard based prototype has a tuning range from 1.2 - 6 GHz. We integrated the SweepSense control and band-select logic into the FPGA onboard the USRP, allowing for direct control from Python or GNURadio. The prototype was able to achieve a maximum speed of 5 ms for the 4.8 GHz bandwidth that it can monitor while being able to extract occupancy measurements for the entire band. This implementation is open source and available at https://github.com/ucsdsysnet/sweepsense.

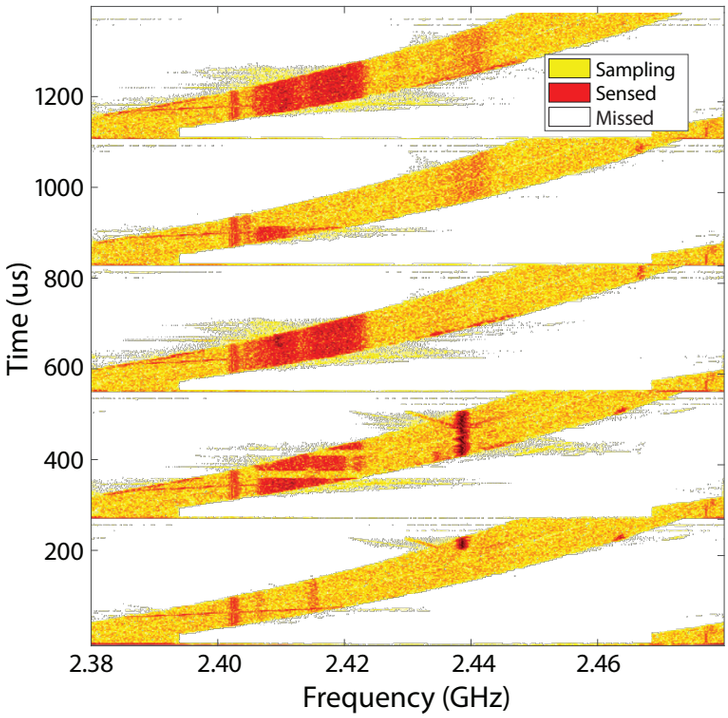

FIGURE 3. Example of 2.4 GHz ISM-band transmissions recovered from SweepSense samples. Wi-Fi OFDM transmissions are visible around 2.41 GHz, a Bluetooth transmission is visible near 2.44 GHz.

To demonstrate SweepSense's capability, we consider two end-to-end, practical case studies in two impact areas. A more detailed treatment of the signal processing used and other benchmarks are presented in the extended paper([9]).

DYNAMIC AND DISTRIBUTED SPECTRUM MANAGEMENT

The SweepSense architecture provides methods to extract useful information (such as protocol) from every transmission sensed. The sensor can be widely deployed (on infrastructure and even vehicles) to allow for real-time monitoring of spectrum utilization (what is transmitted where). Together, this can enable providers to orchestrate complex networks that can even scale to upcoming coexistence-based schemes. Further, having such a sensor deployment could allow reporting of spectrum usage compliance reliably. For example, protocols can be classified and spurious transmissions shut down before they interfere with incumbent systems.

Protocol Classification in Crowded Bands: In scenarios where different users and devices access the same piece of spectrum, it is advantageous to understand what was transmitted and when. The first step in this direction is to detect that a signal exists and map it to a well-known protocol. We pick the ISM band at 2.4 GHz since it hosts a wide variety of protocols (Figure 3). For our evaluations, we choose signals of four protocols (two wideband and two narrowband) representative of this band. The system is evaluated at different sweep rates (frequency swept per unit time) for its accuracy of detection and classification of protocols.

We show that SweepSense is correctly able to classify all types of signals with more than 95% accuracy at decodable SNRs of the order of 10 dB at the fastest tested sweep rate of 125 us per 100 MHz. The system retains this accuracy for sub-noise floor SNRs for some of the signals due to the noise averaging properties of cyclostationary signal processing.

Radar Detection in Shared Bands: Recall that highly reliable sensors that monitor the entire spectrum for incumbent communications are arguably one of the most critical parts enabling CBRS adoption. We show that SweepSense can be configured to sweep twice the FCC required bandwidth for CBRS while being able to detect pulse-type radars within specifications, demonstrating excellent potential coexistence based applications.

CYBER-PHYSICAL DATA ASSOCIATION AND SENSING

Wireless signals could be considered as an abstract representation of underlying real-world phenomena. For example, a spike in utilization around 130 MHz (VHF Airband) corresponds to a busy time at the local airport. At the same time, a flurry of Bluetooth signals at a particular location could be due to a wireless headset nearby. Indeed, the wireless spectrum around us is a rich source of Big Data that has remained untapped. Being able to sense a wide bandwidth quickly, SweepSense opens up a way to gather this data tractably and allow for novel applications to be built on top of it. Distributed deployments further add a spatial dimension to this data source. We believe that scenarios in which wireless signals like LTE are used to predict event-based information and human behavior are within reach with this platform.

Fine-grained LTE Usage Statistics: By configuring SweepSense to monitor multiple LTE bands simultaneously, fine-grained Downlink (DL) usage statistics can be computed based on the power in each channel. We show simultaneous usage statistics for five LTE base stations of different bandwidths (totaling 75 MHz) scattered between the 1.9 - 2.1 GHz licensed spectrum. Further, SweepSense can extract these statistics at a slot level granularity. Interestingly, while performing these experiments, we noted the correlation between channel usage and day-to-day occurrences. Notably, there was a dip in traffic during lunchtime, and spikes during special events with many in attendance.

CONCLUDING REMARKS

SweepSense represents a fundamental change in the approach to spectrum sensing. Sweeping a receiver's center frequency while simultaneously sampling can achieve excellent time resolution sensing across a broad bandwidth. We show that SweepSense can be implemented with simple modifications to off-the-shelf radio architectures. The system shows great potential in transforming impact areas like spectrum management and data-driven sensing applications. ■

Yeswanth Guddeti is a PhD student from UCSD. He received his bachelor's degree from the Indian Institute of Technology, Madras. His research interest is in developing radio systems that explore new paradigms to interact with the RF spectrum for applications, such as dynamic spectrum sharing, device-agnostic localization, and Cyber-Physical Data Association & Sensing. He received the Qualcomm Innovation Fellowship 2019.

Raghav Subbaraman is a PhD student in the Electrical and Computer Engineering Department at UC San Diego. His research interests broadly lie in the domain of wireless sensing, with a special focus on sensor data integration to improve communication infrastructure. Previously, he completed his BTech and worked with the Indian 5G Testbed Project at IIT Madras.

Moein Khazraee is a PhD candidate in Computer Science and Engineering Department at UC San Diego. His research interests include hardware specialization for networking systems, frameworks to improve software and hardware interface, software-defined radios, lowering hardware systems NRE costs, and ASIC Clouds. Khazraee received an MS in computer science from the University of California, San Diego.

Aaron Schulman is an Assistant Professor at the University of California, San Diego. He develops and deploys measurement tools to study the efficiency, reliability, and security of the mobile infrastructure. He earned his PhD in Computer Science from the University of Maryland where he studied Internet reliability, and he was a Postdoc at Stanford where he found bottlenecks in cellular infrastructure and founded a company that helped Google improve the battery life of the Chrome web browser. He received the 2013 ACM SIGCOMM Doctoral Dissertation Award.

Dinesh Bharadia is an Assistant Professor at UCSD. His research interests lie in the design and development of systems that advance the theory and implementation of modern wireless communication, sensing, and networking systems. He worked at Kumu Networks to commercialize his PhD thesis research on full-duplex radio and was also awarded Forbes 30under30, Marconi Young Scholar Award, and the MIT TR35 top 35 Innovators Under 35 awards for the same.

REFERENCES

[1] U.S. Government. CFR title 47 section 96.67 Environmental Sensing Capability.

[2] A. P. Goodson. A multi-function, broad band, high dynamic range RF receiver. Technical report, OneRadio, 2017.

[3] Anritsu. MS2840A Spectrum Analyzer. https://www.anritsu.com/en-us/test-measurement/products/ms2840a

[4] Roundup of Software Defined Radios- rtl-sdr.com [online] Available: https://www.rtl-sdr.com/roundup-software-defined-radios/

[5] N200/N210 - Ettus Knowledge Base, 2020, [Online] Available: https://kb.ettus.com/N200/N210

[6] L. Shi, P. Bahl, and D. Katabi. 2015. Beyond sensing: Multi-GHz realtime spectrum analytics. In *Proc. Symposium on Networked Systems Design and Implementation (NSDI).*

[7] H. Hassanieh, L. Shi, O. Abari, E. Hamed, and D. Katabi. 2014. GHz-Wide sensing and decoding using the sparse fourier transform. In *Proc. IEEE Conference on Computer Communications (INFOCOM).*

[8] W.A. Gardner, A. Napolitano, and L. Paura (2006). Cyclostationarity: Half a century of research. *Signal Processing*, 86(4), 639-697.

[9] Y. Guddeti, R. Subbaraman, M. Khazraee, A. Schulman, and D. Bharadia. Sweepsense: Sensing 5 ghz in 5 milliseconds with low-cost radios. 2019. *Proceedings of the 16th USENIX Symposium on Networked Systems Design and Implementation (NSDI '19)*, 317–330.