DRAFT CBRS Ecosystem Sharing Assessment (SEA) Test and Metrology Test Plan

Duncan A. McGillivray, Todd Schumann, Thao T. Nguyen
Douglas Boulware, M. Keith Forsyth
Dazhen Gu, C. Lee Joyce
Elyssa Kaplan, Mark Krangle
Daniel G. Kuester, Joseph R. Mruk
Anthony W. Romaniello, Aric W. Sanders
Robert Scheeler, Stephen Segro

NASCTN Program Manager: Melissa M. Midzor
NASCTN Project Manager: Keith A. Hartley

February 23, 2023

NASCTN

National Advanced Spectrum and Communications Test Network

DISCLAIMER

Certain commercial entities, equipment, or materials may be identified in this document in order to describe an experimental procedure or concept adequately. Such identification does not imply any recommendation or endorsement by National Institute of Standards and Technology (NIST) or National Advanced Spectrum and Communications Test Network (NASCTN), nor is it intended to imply that the materials or equipment identified are necessarily the best available for the purpose.

This document is available free of charge from:
https://www.nist.gov/programs-projects/cbrs-sharing-ecosystem-assessment

Preface

The National Advanced Spectrum and Communications Test Network (NASCTN) provides an impartial, scientifically rigorous forum for addressing spectrum-sharing challenges in an effort to accelerate the deployment of wireless technologies among commercial and federal users and to measure the impacts of spectrum dependent systems deployments.

NASCTNs mission is to provide robust test processes, validated measurements data, and statistical analysis necessary to develop, evaluate, and deploy spectrum sharing technologies that can incrase access to the spectrum by both federal agencies and non-federal spectrum users.

Representatives from Defense Information Systems Agency (DISA) Defense Spectrum Organization (DSO) submitted a proposal to NASCTN to measure emissions in the Citizens Broadband Radio Service (CBRS) band to collect data required for DISA DSO to ascertain the effectiveness of the sharing ecosystem between CBRS systems as managed by Spectrum Access Systems (SASs), and Department of Defense (DoD) systems as monitored by Environmental Sensing Capabilitys (ESCs). The collect data will be used by DoD to provide insight into the sharing ecosystem's effectiveness, and track changes in the spectrum environment over time.

After the NASCTN Steering Committee accepted the DSO proposal as a NASCTN test, NASCTN assembled a test plan development team comprised of experienced engineers and other professionals. The results of that work plus the knowledge and creativity of the team members culminated in the production of this test plan.

The test plan is designed to yield reproducible measurements. This NASCTN effort focuses on the viability of spectrum sharing in the CBRS band in Always-On Dynamic Protection Areas (DPAs). NASCTN will solicit comments on this test plan from CBRS federal and non-federal stakeholder community.

A description of NASCTN and listing of Charter Members at the time of publication of this report follows.

National Advanced Spectrum and Communications Test Network (NASCTN)

The mission of the NASCTN is to provide, through its members, a network for robust test processes and validated measurement data necessary to develop, evaluate and deploy spectrum sharing technologies that can improve access to the spectrum by both federal agencies and non-federal spectrum users.

NASCTN is a member organization under a charter agreement. Members

- Facilitate and coordinate work with federal, academic, and industry spectrum users to rapidly and cooperatively facilitate spectrum sharing and co-existence studies;
- Work as a partnership to address the interests and equities of all spectrum stakeholders in a fair, equitable, and non-preferential manner; and
- Through sharing of technical resources, with consideration for cost, provide liaison and support to coordinate and leverage existing national capabilities supporting government, academic, and industry testing and evaluation known to improve and expedite spectrum sharing and coexistence.

Charter members at the time of publication of this report are (in alphabetical order):

- Department of Defense Chief Information Officer (DoD CIO)
- National Aeronautics and Space Administration (NASA)
- National Institute of Standards and Technology (NIST)
- National Oceanic and Atmospheric Administration (NOAA)
- National Science Foundation (NSF)
- National Telecommunications and Information Administration (NTIA)

National Institute of Standards and Technology (NIST) hosts the NASCTN capability at the Department of Commerce Boulder Laboratories in Boulder, Colorado.

Technical Contributors

Contributor	Organization	Primary Contribution Areas
Jordan L. Bernhardt	NIST NASCTN	Data edge compute and analysis
M. Keith Forsyth	NIST NASCTN	Sensor system, and fielding
Duncan A. McGillivray	NIST NASCTN	Project technical leadership
Aric W. Sanders	NIST NASCTN	System verification, bench acceptance testing, and calibration
Stephen Segro	NIST NASCTN	Sensor system co-lead, site survey, and fielding
Dazhen Gu	NIST CTL	System verification, bench acceptance testing and calibration
Daniel G. Kuester	NIST CTL	Date edge compute and analysis
John M. Ladbury	NIST CTL	Consultation to sensor system and calibration
Thao T. Nguyen	NIST CTL	Project technical leadership, modeling and simulation
Robert Ballard	NTIA ITS	SCOS software and implementation
Douglas Boulware	NTIA ITS	SCOS software and implementation
Joel Dumke	NTIA ITS	Consultation to modeling and simulation
Bradley Eales	NTIA ITS	SCOS software and implementation
Justin Haze	NTIA ITS	SCOS software and implementation
Peter Mathys	NTIA ITS	Date edge compute and analysis
Anthony W. Romaniello	NTIA ITS	Integrated product, site survey, and sensor fielding
Todd Schumann	NTIA ITS	Project technical leadership, Integrated product
C. Lee Joyce	NASA LaRC / AMA	Site survey, and sensor fielding
Keith Hartley	MITRE	Project management
Mark Krangle	MITRE	Sensor system, and sensor fielding
Elyssa Kaplan	MITRE	Sensor system
Joseph R. Mruk	MITRE	Sensor system
Robert Scheeler	MITRE	Sensor system, and modeling and simulation

Acknowledgments

We would like to thank the many additional individuals and organizations that made valuable contributions over the course of this test plan development phase.

The insights gathered from the National Advanced Spectrum and Communications Test Network (NASCTN) test plan sponsor Defense Information Systems Agency (DISA) Defense Spectrum Organization (DSO), Federal and industry representatives of the Interagency-Joint Working Group on Citizens Broadband Radio Service (CBRS) were crucial to the framing of the development effort. We thank industry for hosting NASCTN at Wireless Innovation Forum (WInnForum), On-Go Alliance, and the On-Go Alliance board members meeting.

In addition, we thank Rebecca Dorch of National Telecommunications and Information Administration (NTIA) Institute for Telecommunication Sciences (ITS) for providing crucial access, context, support and insights.

We thank United States Marine Corps (USMC) representatives Seth Nagy, and Chief Warrant Officer Figueroa for their coordination and site survey efforts at USMC Camp Pendleton, which were formative to this test plan.

The initial NASCTN Sharing Ecosystem Assessment (SEA) sensor demonstration would not have been possible without the explicit support of the 3.5 GHz Transition Plan through Department of Defense (DoD) Chief Information Officer (CIO), DISA DSO, and Department of Navy (DoN) CIO. We recognize the consultation and technical engagement by Naval Sea Warfare Center Dahlgren Division's Stacy Barker, Andrew Simpson, Han Lee, and Mike Mearns.

For the hosting, coordination, and deployment of NASCTN SEA sensors we thank Professor Bill Moore of Hampton University, Kenneth Dudley of the Longbow Group, and Randy Plotkin of the Port of Virginia.

The technical team appreciates the opportunity and continued encouragement to participate in this project and recognizes their leadership: NASA LaRC Branch head Jay Ely, and Test Range Manager Charmaine Franck
NIST CTL Division Chiefs Paul Hale, Melissa Midzor, Richard Rouil
NTIA ITS Division Chiefs Michael Cotton, Julie Kub, and Charles Dietlein

Also, the technical team thanks the NASCTN program manager, Melissa Midzor, the NASCTN administrative officer Amanda Hyman, division office manager Julie Zumhofe, and former NASCTN project manager Matthew Briel for their efforts.

Finally, thanks to Kasey Pugh of DISA DSO and supporting staff led by Tim MUll of The MITRE Corporation (MITRE) for their sponsorship, continued guidance and feedback throughout the evolution of the SEA project.

Executive Summary

The purpose of this National Advanced Spectrum and Communications Test Network (NASCTN) project is to provide data-driven insight into the Citizens Broadband Radio Service (CBRS) sharing ecosystem's effectiveness between commercial and Department of Defense (DoD) incumbent systems, and to track changes in the spectrum environment over time. CBRS is the "first of a kind" nationwide shared spectrum ecosystem in the 3550-3700 MHz band. It includes a 3 tier approach (Incumbent, Priority Access License (PAL), General Authorized Access (GAA)), specific sharing stipulations as notated in the Federal Communications Commission (FCC) Part 96 rules, and Environmental Sensing Capability (ESC) which detects incumbents and Spectrum Access System (SAS) that coordinate spectrum access. This project is part of the DoD 3.5 GHz Transition Plans that focus on ecosystem validations, environmental assessments, and continued engagement on refining the CBRS infrastructure.

The FCC established the CBRS in the 3550-3700 MHz band, and created a first of its kind three-tiered access and authorization framework to accommodate Federal incumbent and non-federal users which are tiered as PAL and GAA users. The FCC auctioned the band (Auction #105) for commercial use in July 2020. A system has been established to enable this sharing, which is composed of two main parts: SASs which perform automated spectrum use coordination, and ESCs that detect incumbent use of the band and transmit that information back to the SASs.

This project is divided into four major components:

- 1 - Passive Observation: Characterize aggregate emissions within the CBRS band 3550-3700 MHz with and without Dynamic Protection Area (DPA) activations, in at least two coastal DPAs. Assess ecosystem performance to timely respond, and measure increase in noise floor due to CBRS deployments over time.
- 2 - Passive Observation: Characterize aggregate emissions within the CBRS band 3550-3700 MHz in at least one Always-On DPA. Assess ecosystem performance to limit CBRS emissions in the Always-On DPA.
- 3 - Active Experimentation: Evaluate ESC detection performance in the Field (with Navy Ship). For a representative set of ESCs in the field, assess true-positive and false-positive rates and independently verify incident power levels at ESC sensor locations.
- 4 - Long Term support - Collection and analysis of Passive Observations over 4 years, and support transfer of data to a DoD data repository.

This test plan addresses Passive Observation in the Always-On DPA. Separate test plans will be developed for Passive Observation in Coastal DPAs and Active Experimentation.

NASCTN projects follow a 5 stage open, transparent, comprehensive process for developing and executing independent, scientifically based test plans, validating test results, and reporting findings. This test plan is part of Stage III, which includes conducting investigations, building prototypes/test beds, design of experiment, and pilot studies. During this stage, the team completed three major efforts to inform the development of this test plan:

- Sensor site selections - Site surveys for multiple locations within several DPAs were conducted. One Always-On DPA was chosen for this test plan.
- Modeling - Incorporated existing methods and industry provided data to better inform site selections, and to identify measurements to achieve project goals.
- Prototype Sharing Ecosystem Assessment (SEA) sensor/ Edge Data collection systems - Provide insights into CBRS ecosystem behavior at representative locations; exercise remote sensing nodes which will perform time autonomous, calibrated measurements with quantified levels of uncertainty; and trial edge-compute data products in situ.

In addition, extensive outreach to the CBRS community was undertaken in order to aid in resolving underlying assumptions of the test framework, gain improvements and insights to arrive at a tractable and scoped test plan, and context to broaden utility of the NASCTN test and data collection to benefit the CBRS community.

The general approach to achieve the objectives of this test plan consists of deploying a quantity of custom-build emission collection sensors and analysis software designed to detect, observe, and collection CBRS emissions. A number of sensors will be strategically located within an Always-On DPA to continuously monitor the CBRS channels for any emissions that exceed the threshold for Always-On DPA. The collection system and software collects basic occupancy data such as maximum, minimum, mean, and median power collected by repeatedly scanning each of the CBRS 10 MHz channels. The emissions are analyzed onsite and collected data is transmitted back to NASCTN for storage and further analysis.

NASCTN will solicit and adjudicate comments on this test plan from the CBRS stakeholders within the Federal and non-federal communities.

Contents

Acronyms

3GPP 3rd-generation partnership project
AC alternating current
ADC analog-to-digital converter
AGL above ground level
AMA Analytical Mechanics Associates
AMSL above mean sea level
APD amplitude probability distribution
API application programming interface
AWGN additive white Gaussian noise
CAC Common Access Card
CBRS Citizens Broadband Radio Service
CBSD Citizens Broadband Radio Service device
CCDF complementary cumulative distribution function
CIO Chief Information Officer
COTS commercial off-the-shelf
CTL Communications Technology Laboratory
CW continuous-wave
DANL displayed average noise level
DFT discrete Fourier transform
DISA Defense Information Systems Agency
DoC Department of Commerce
DoD Department of Defense
DoN Department of Navy
DPA Dynamic Protection Area
DSO Defense Spectrum Organization
DSP digital signal processing
DTFT discrete-time Fourier transform
E-DPA ESC-monitored DPA
EIRP effective isotropic radiated power
ENR excess noise ratio
ENBW equivalent noise bandwidth
ENR excess noise ratio
ESC Environmental Sensing Capability
FAD full activity dump
FCC Federal Communications Commission
FFT fast Fourier transform
FIR finite impulse-response
GAA General Authorized Access
GB-DPA ground-based DPA
GEOID geographic identifier
GMM Green Mountain Mesa
GUI graphical user interface
HU Hampton University
IEEE Institute of Electrical and Electronics Engineers
IIR infinite impulse-response
IP internet protocol
IQ in-phase and quadrature
ITM Irregular Terrain Model
ITS Institute for Telecommunication Sciences
JSON JavaScript object notation
KML keyhole markup language
LAN local area network
LaRC Langley Research Center
LNA low noise amplifier
LTE long-term evolution
LZMA Lempel-Ziv-Markov chain algorithm
MITRE The MITRE Corporation
Msps mega samples per second
MUF microwave uncertainty framework
NASA National Aeronautics and Space Administration
NASCTN National Advanced Spectrum and Communications Test Network
NIT Norfolk International Terminal
ND noise diode
NEMA National Electrical Manufacturer's Association
NF noise figure
NIST National Institute of Standards and Technology
NPS National Park Service
NTIA National Telecommunications and Information Administration

OSM Office of Spectrum Management
PAL Priority Access License
P1dB 1 dB gain compression
PFP periodic frame power
PIV Personal Identity Verification
PSD power spectral density
RF radio frequency
RMS root mean square
RSA real-time spectrum analyzer
SAS Spectrum Access System
SCOS spectrum-characterization and occupancy sensing
SEA Sharing Ecosystem Assessment
SigMF signal metadata format
SNR signal to noise ratio
SPDT single pole double throw
SPU signal processing unit
SQL Structured Query Language
STIG Security Technical Implementation Guide
SWaP size, weight and power
TDD time division duplex
UPS uninterruptible power supply
USFWS United States Fish and Wildlife Service
USN United States Navy
USB universal serial bus
USMC United States Marine Corps
VPA Virginia Port Authority
WGS-84 World Geodetic System 1984
WInnForum Wireless Innovation Forum

List of Figures

List of Tables

Chapter 1

Introduction

Defense Information Systems Agency (DISA) Defense Spectrum Organization (DSO) proposed a measurement campaign to the National Advanced Spectrum and Communications Test Network (NASCTN) to collect Citizens Broadband Radio Service (CBRS) emissions in the 3550-3700 MHz band to allow Department of Defense (DoD) to collect data required for DISA DSO to ascertain the effectiveness of the sharing ecosystem between CBRS systems as managed by Spectrum Access Systems (SASs), and DoD systems as monitored by Environmental Sensing Capabilitys (ESCs). The collected data provides DISA DSO insight into the sharing ecosystem's effectiveness, and track changes in the spectrum environment over time.

This test plan is focused on the technical approach to collect CBRS emissions in Always-On Dynamic Protection Areas (DPAs) through a series of deployed sensors and software to collect and report CBRS emissions for visualization and analysis. The test plan describes the methodology, sensor and software design, data collection, modeling, and data visualisation/analysis to provide yearly reports and accompanying data on the collected CBRS emissions.

1.1 Background

1.1.1 NASCTN Process

NASCTN projects follow a 5-Stage open, transparent, comprehensive process for developing and executing independent, scientifically based test plans, validating test results, and reporting findings as shown in Figure 1.1. This serves as a common architecture across all NASCTN projects.

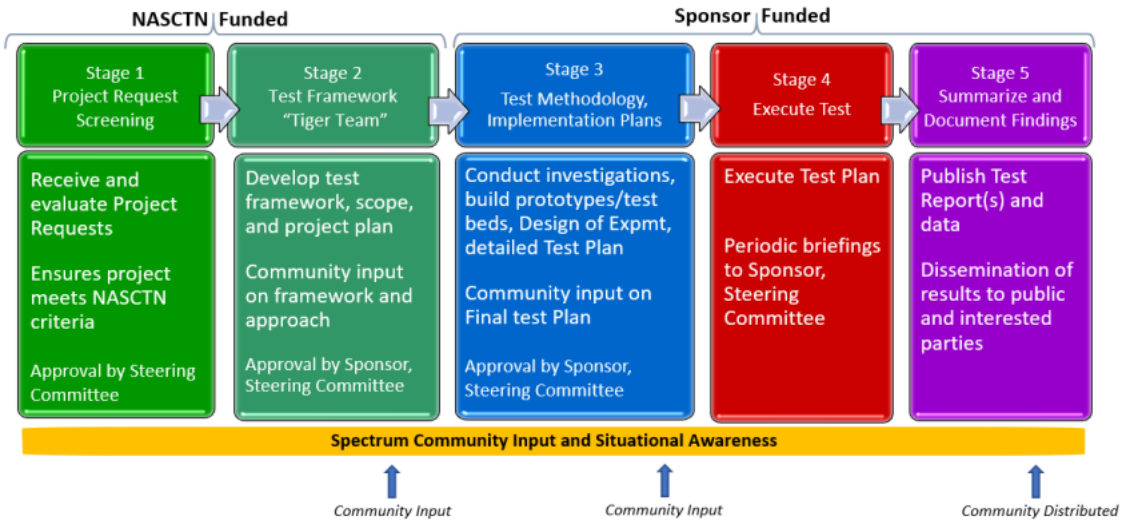

Figure 1.1: NASCTN's 5 Stage Process

1.1.2 CBRS Ecosystem

The Federal Communications Commission (FCC) established the CBRS in the 3550-3700 MHz band, and created a first of its kind three-tiered access and authorization framework to accommodate Federal incumbent and non-federal users which are tiered as Priority Access License (PAL) and General Authorized Access (GAA) users as shown in Figure 1.2. The FCC auctioned the band (Auction #105) for commercial use in July 2020. A system has been established to enable this sharing, which is composed of two main parts: SASs which perform automated spectrum use coordination, and ESCs that detect incumbent use of the band and transmit that information back to the SASs.

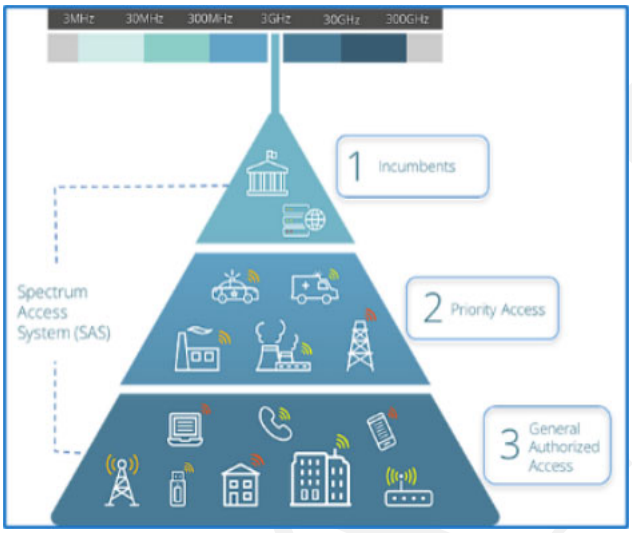

Figure 1.2: FCC's three-tiered access for CBRS. Reference: https://ongoalliance.org/the-technology-behind-spectrum-sharing-the-spectrum-access-system/

1.1.3 Test Request

DoD 3.5 GHz Transition Plans focus on ecosystem validations, environmental assessments, and continued engagement on refining the CBRS infrastructure. One component was identified for submission to NASCTN: to evaluate the effectiveness of the CBRS sharing ecosystem that manages the co-existence between Federal incumbent systems and commercial users, via an independent trusted agent.

The DISA DSO submitted a CBRS Sharing Ecosystem Assessment (SEA) test request to NASCTN to achieve four objectives regarding the CBRS band:

- Provide data driven insights into:
  1. Efficacy of permanent sharing between CBRS systems as managed by SASs and ESCs
  2. Noise floor measurements through continuous automated observations
  3. CBRS emissions within the Always-On DPAs at Army and Marine sites
- Support
  1. The development and management of a measurement data repository

This DISA DSO test request was divided into four major components:

1. Passive Observation: Characterize aggregate emissions within the CBRS band 3550-3700 MHz with and without DPA activations, in at least two coastal DPAs. Assess ecosystem performance to timely respond, and measure increase in background emissions due to wireless system deployments over time.

2. Passive Observation: Characterize aggregate emissions within the CBRS band 3550-3700 MHz in at least one Always-On DPA. Assess ecosystem performance to limit CBRS emissions in the Always-On DPA.
3. Active Experimentation: Evaluate ESC detection performance in the Field (with Navy Ship). For a representative set of ESCs in the field, assess true-positive and false-positive rates and independently verify incident power levels at ESC sensor locations.
4. Long Term support - Collection and analysis of Passive Observations over 4 years, and support transfer of data to a DoD data repository.

The project is intended to be conducted over 5 years and is comprised of two phases:

- Phase 1: Establish the measurement system requirements; develop, build, deploy a metrology sensor system to measure the CBRS ecosystem performance at DoD sites (e.g., Camp Pendleton, CA) and potentially other non-DoD sites.
- Phase 2: Long-term analysis phase, with continuous emission level measurements and annual deep dive measurements and analysis.

1.2 Objective

The overall objective of this test plan to develop a methodology to leverages field-collected emissions in the CBRS band in at least one Always-On DPA combined with modeling to provide the data needed for the DoD to evaluate the effectiveness of spectrum sharing in the 3550-3700 MHz band.

A yearly report will be published for CBRS stakeholders with the results of the collected data, modeling, and analysis.

1.3 Scope

The test method discussed in this plan focuses on field-captured emissions in the CBRS band and complimented by modeling and simulation. This test plan is designed to collect data to assist Federal stakeholders to independently confirm effective spectrum sharing. The test is not designed to ascertain individual Citizens Broadband Radio Service device (CBSD), SAS Administrators, or ESC operators compliance with established FCC rules for the CBRS band.

1.4 Deliverables

The following lists the deliverables of this effort:

- A test plan which includes a repeatable test methodology to collect CBRS ecosystem data to inform on the emission levels within the CBRS band for Always-On DPAs.
- A set of validated data with uncertainties from field collections and modeling to augment the interpretation of field collects.
- An initial set of analysis and visualization tools to view and interpret the collected data
- A yearly report summarizing the data collected over the preceding year
- Support establishing a repository for stakeholder access to collected, processed, and analyzed data
- Work to identify mechanisms and approaches to share interim working-level data on a quarterly basis

Chapter 2

Modeling and Simulation

Modeling and simulation are utilized to provide insight into the Citizens Broadband Radio Service (CBRS) ecosystem and to augment operational test design and evaluation of the National Advanced Spectrum and Communications Test Network (NASCTN) CBRS Sharing Ecosystem Assessment (SEA) project.

2.1 Models

This section describes the CBRS deployment model, propagation model, and Dynamic Protection Area (DPA) move list algorithm used for incumbent protection.

2.1.1 Obfuscated Data on CBRS Deployment

The Spectrum Access System (SAS) Administrators provided the National Telecommunications and Information Administration (NTIA) Institute for Telecommunication Sciences (ITS) obfuscated data on CBRS deployment and NASCTN received obfuscated full activity dump (FAD) data for three selected areas. These include two areas near DPA West-14 and DPA East-1, and one area encompassing Boulder County and Broomfield County in Colorado. For areas near DPA West-14 and DPA East-1, the obfuscated data were filtered out using the respective Category B Citizens Broadband Radio Service device (CBSD) neighborhood distance plus a 100 km distance buffer as shown in Figure 2.1a and Figure 2.1b. Detailed information related to DPAs, e.g., DPA-specific neighborhood distances and protection criteria, can be found in DPA keyhole markup language (KML) files provided by NTIA [2]. Since there is no DPA near Boulder, the geographical boundaries of Boulder County and Broomfield County were used to filter out the obfuscated data as shown in Figure 2.1c.

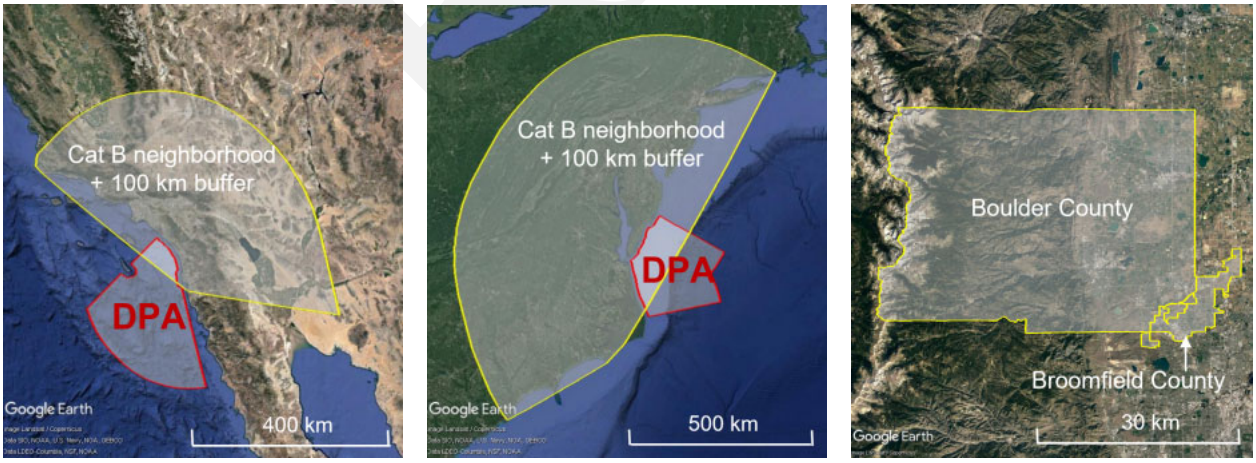

(a) DPA West-14 (b) DPA East-1 (c) Boulder and Broomfield Counties

Figure 2.1: Geographical masks used to filter out obfuscated data for three areas of interest.

The procedure used to generate obfuscated FAD data can be summarized as follows:

1. For each CBSD "registration" in the FAD data, "radioTechnology" within "airInterface" was obfuscated to replace any proprietary protocols. If "airInterface" was not "E_UTRA", "NR", or "CAMBIUM", it was replaced with "OTHERS". In addition, "latitude" and "longitude" were removed from the "installationParam", and a geographic identifier (GEOID) of the census block containing the latitude and longitude was added.
2. The FAD data was then aggregated into distributions for each census block. Under each census block GEOID, there are three separate distributions for "Category_A", "Category_B", and "Unknown_Category" CBSDs. Further, each category contains distributions of:
   - "airInterface": contains the number of CBSDs that have "E_UTRA", "NR", "CAMBIUM", or "OTHERS" air interfaces.
   - "heightAGL", "heightAMSL", and "unknownHeight": each height distribution contains the number of CBSDs that have above ground level (AGL), above mean sea level (AMSL), or unknown heights such that $h_i \leq h_{CBSD} < h_{i+1}$, where h_{CBSD} is the CBSD height, h_i and h_{i+1} are two consecutive bin heights with 3 m bin size.
   - "maxEirp": contains the number of CBSD grants that have max effective isotropic radiated power (EIRP) values such that $\text{maxEirp}_i < \text{maxEirp}_{CBSD} \leq \text{maxEirp}_{i+1}$, where maxEirp_{CBSD} is the maxEirp value of the grant, maxEirp_i and maxEirp_{i+1} are two consecutive bin values. The bins are every 5 dB down from 37 dBm except the lowest bin which includes values of $0, 1, 2$ dBm.
   - "channel_use": contains the number of CBSD grants operating on each 10 MHz channel from 3550 MHz to 3700 MHz.
   - "bandwidth": contains the number of CBSD grants that received each specified grant bandwidth.
   - "count": contains a count of the total number of CBSDs.
   - "indoor_deployment_count": contains a count of the total number of indoor CBSDs.

The obfuscated data were used to generate randomized CBSD deployments for each area. Two categories of CBSDs, i.e., low-power indoor Category A CBSDs and high-power outdoor Category B CBSDs, were generated separately for each census block GEOID provided in the obfuscated data file. Note that "Unknown_Category" CBSDs have not been populated in the latest obfuscated data. And for each category, the CBSD registration request data and grant request data were randomly generated from the obfuscated distributions, and then the grants were randomly assigned to the CBSDs. Note that in the CBRS standards [3], a "grant" is an authorization provided by a SAS to a CBSD, subject to a Heartbeat exchange, to transmit using specified operating parameters; and a CBSD may have multiple grants.

2.1.2 Propagation Model

The Wireless Innovation Forum (WInnForum) CBRS standards for SAS general requirements specify in [4, R2-SGN-03] that a SAS shall use the NTIA ITS Irregular Terrain Model (ITM) [5] (also known as the Longley-Rice model) in point-to-point mode for use in DPA protection for Initial Certification. The ITM model is a terrain-based model and it does not include clutter loss. However, to account for building attenuation, which is also absent in the ITM model, additional loss of 15 dB is ascribed if the CBSD is located indoors [4, R2-SGN-03]. Other parameters used in the ITM model are provided in another SAS general requirement [4, R2-SGN-17].

An open-source reference implementation of the ITM model is available at the WInnForum SAS Testing and Interoperability repository [6]. Terrain and other data used by the model are available at the Common WInnForum Data Libs [7].

2.1.3 WInnForum Reference DPA Move List Algorithm

Regulatory rules [8] require that the SAS administrators manage their CBSD transmissions to protect the operations of existing incumbents in the band. To meet the requirement, the WInnForum CBRS standards specify the DPA protection procedure, also known as the move list algorithm, to be executed by all SASs. Given a set of CBSD transmissions having or requesting grants that overlap in frequency with a protected frequency range, the move list algorithm identifies which transmissions must be suspended (and possibly relocated to a different channel) to avoid

excessive interference in a protected federal incumbent area.

Details of the algorithm and its reference implementation can be found in [4, R2-SGN-24] and [6]. In general, for any protected frequency range ch and for any protection point p within a given DPA, the move list algorithm first determines a set of N_c CBSDs having or requesting a grant that includes any portion of the protected frequency range ch, and that are within a neighborhood of the protection point p. It then computes the path loss from each grant to the protection point p using the ITM model described in Section 2.1.2. The algorithm applies a reference receive antenna pattern for protection points with a main lobe beamwidth of 3 degrees, a reference gain of 0 dBi within the beamwidth, and a reference gain of -25 dBi outside of main lobe [4, R2-IPM-04]. The algorithm then chooses a subset of the grants that must be suspended (relocated) such that the 95^{th} percentile of the aggregate emission, from N_c grants to the protection point p with a minimum of 2000 Monte Carlo trials [4, R2-IPM-03], is below a predefined threshold. The grants that must be removed to meet the protection threshold are placed on the move list. This process is repeated for every protection point in the protection area DPA. The move list for the entire protection area DPA is the union of the move lists of the protection points inside the DPA.

The primary output of the move list algorithm is the move list. For the purpose of this modeling and simulation, a neighbor list and a keep list for the protection area DPA are also considered. The neighbor list is defined as the union of the neighbor lists of the protection points inside the DPA. Whereas, the keep list is defined as a complement of the move list, and it contains grants that may remain active on the protected channel when the DPA is active.

2.2 Simulation

This section showcases examples of DPA move list simulation results applicable to Task 1 and Task 2. From the obfuscated data, randomized CBSD grant deployments can be generated for each 10 MHz channel in the 3550-3650 MHz band and input to the reference implementation of the DPA move list algorithm. The outputs include the neighbor list, move list, and keep list grants for each 10 MHz protected channel for a DPA under consideration. Examples of simulation results for Task 1 and Task 2 are shown in the following subsections.

2.2.1 Examples of Task 1 Simulation Results

The two DPA West-14 and DPA East-1 considered in Task 1 are monitored by Environmental Sensing Capability (ESC) systems. The ESC notifies SASs to activate one or more DPAs when it detects incumbent signals potentially inside any of those DPAs. This triggers SASs to suspend or relocate CBSD grants on the move list.

Figure 2.2 and Figure 2.3 show examples of grants on the neighbor list, move list, and keep list for DPA West-14 and DPA East-1, respectively. The number of grants per 10 MHz channel are in the hundreds for DPA West-14 and in the thousands for DPA East-1.

Note that, due to proprietary concerns, the CBSD grant locations in these figures are used for illustration purposes and they are not locations derived from the obfuscated data. Specifically, these CBSD grant locations are derived from the full CBSD deployment model which was used for the 3.5 GHz exclusion zone analyses by the NTIA Office of Spectrum Management (OSM) in [9].

2.2.2 Examples of Task 2 Simulation Results

For Task 2, the always-on DPA at Camp Pendleton is considered. The DPA is used to protect incumbent systems operating below 3500 MHz. The protection criteria, e.g., protection thresholds and neighborhood distances, vary by 10 MHz channels in the 3550-3650 MHz band[2]. In this simulation, a default method, implemented by the WInnForum [6] with example values of $(150, 0, 50, 40, 0.5, 0, 0.5, 0)$, is used to get the protection points inside the DPA for move list calculation. Since the DPA is always-on, the grants on the move list are automatically suspended until they come out of the move list in the next move list calculation.

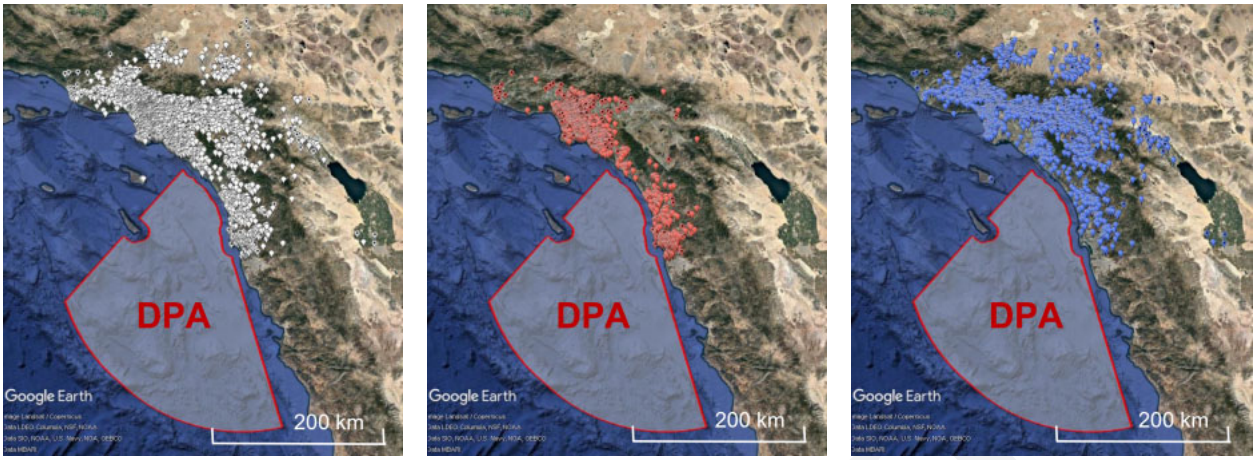

(a) Example of neighbor list grants (b) Example of move list grants (c) Example of keep list grants

Figure 2.2: Examples of simulation results for DPA West-14 in San Diego, California.

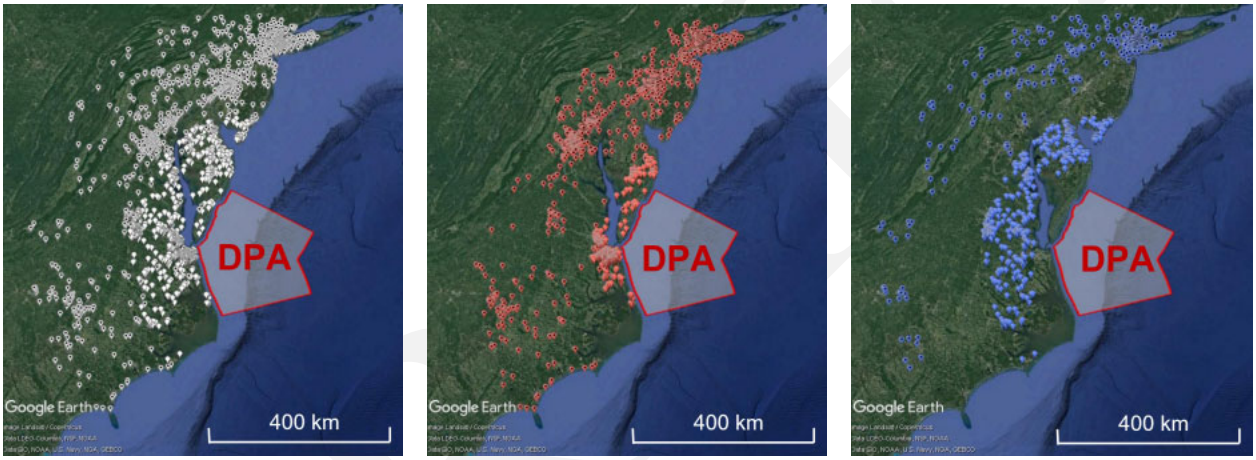

(a) Example of neighbor list grants (b) Example of move list grants (c) Example of keep list grants

Figure 2.3: Examples of simulation results for DPA East-1 in Norfolk, Virginia.

675 Figure 2.4 shows the examples of grants on the neighbor list, move list, and keep list near the always-on DPA Camp Pendleton. There are not many CBSD grants in the neighborhood of the DPA Camp Pendleton. Note that the CBSD grant locations shown in the figures are only for illustration purposes due to proprietary concerns.

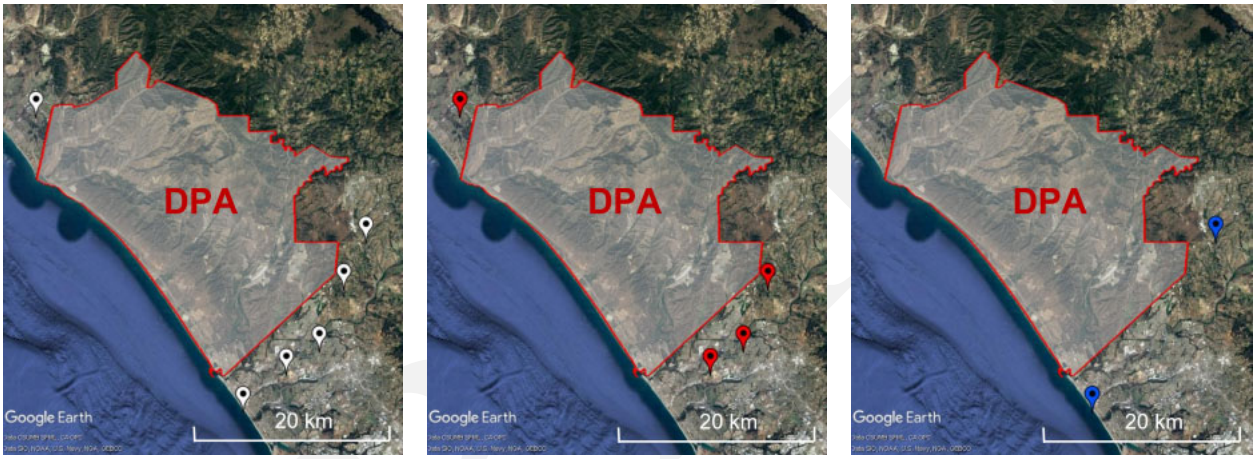

(a) Example of neighbor list grants (b) Example of move list grants (c) Example of keep list grants

Figure 2.4: Examples of simulation results for always-on DPA Camp Pendleton.

Chapter 3

Modeling Informed Site Selection

To assess the effectiveness of spectrum sharing in the Citizens Broadband Radio Service (CBRS) ecosystem, the National Advanced Spectrum and Communications Test Network (NASCTN) Sharing Ecosystem Assessment (SEA) project will need to deploy multiple sensors in the vicinity of Dynamic Protection Areas (DPAs). Given the vast geographical areas of DPAs, modeling is used as a tool to help informing site selections for SEA sensors to measure emissions in the CBRS band.

3.1 Methodology Overview

This section describes the aggregate emissions model used to calculate Citizens Broadband Radio Service device (CBSD) emissions at a given SEA sensor location. An example of the calculation is detailed subsequently in Section 3.1.2 for a prototype sensor deployed at Green Mountain Mesa (GMM) in Boulder, Colorado.

3.1.1 Applied Aggregate Emission Model

Given a SEA sensor location s and any 10 MHz channel ch in the 3550-3700 MHz band, the aggregate emission model first determines a set of N grants that are within the neighborhood or vicinity of the sensor s and requesting or operating on any portion of the channel ch. The model then computes the emission contribution $I_{i\ (\text{dBm})}$ from each individual grant i, for $1 \leq i \leq N$, to the sensor s. Finally, it sums in linear power emission contributions from all N grants together to obtain the aggregate emission received at the sensor.

When computing the emission contribution $I_{i\ (\text{dBm})}$ of an individual grant i, the model uses a simple link budget equation as follows:

$$I_{i\ (\text{dBm})} = EIRP_i - L_i + G_{s,i} \tag{3.1}$$

where $EIRP_i$ is the effective isotropic radiated power (EIRP) value of the grant i within the channel ch (dBm), L_i is the path loss from the grant i to sensor s (dB), and $G_{s,i}$ is the 2D antenna gain of the sensor s in the direction of the grant i (dBi).

Let's define $\{I_{1\ (\text{dBm})}, \cdots, I_{i\ (\text{dBm})}, \cdots, I_{N\ (\text{dBm})}\}$ as a set of N independent, but not necessarily identical, random variables, each representing the emission contribution from a grant i to sensor s on channel ch. The associated emission contribution in linear scale (mW) can be computed as $I_i = 10^{I_{i\ (\text{dBm})}/10}$. Consequently, the aggregate emission I (mW) of N grants is the sum of the emission contribution of grants $I_1, \cdots, I_i, \cdots, I_n$ as follows

$$I = \sum_{i=1}^{N} I_i = \sum_{i=1}^{N} 10^{I_{i\ (\text{dBm})}/10} \tag{3.2}$$

and I (mW) can be converted to log scale by $I_{\ (\text{dBm})} = 10 \log_{10} I$ (dBm).

3.1.2 Example of Aggregate Emission Calculation

To further elaborate on the aggregate emission calculation, an example of such calculation is presented for a small area in Boulder, Colorado. Since there are measurements collected at a prototype SEA sensor in the area (see Section 8.3.1

for more details), it's beneficial to compare the aggregate emission calculated using obfuscated full activity dump (FAD) data with the measured power per channel in the 3550-3700 MHz band. Furthermore, Monte Carlo simulations are performed on different CBSD deployments to determine if the random deployments from the same obfuscated data produce significantly different results.

3.1.2.1 Randomized CBSD Deployment and Sensor Location

A randomized CBSD deployment was generated from the obfuscated FAD data for Boulder County and Broomfield County in Colorado. The obfuscated FAD data was provided by the Spectrum Access System (SAS) Administrators in January, 2023. More Category B CBSDs than Category A CBSDs have been deployed in the area. In total, there are more than 250 CBSDs having approximately 800 grants and operating on 15 channels in the 3550-3700 MHz band. Figure 3.1 shows an example of CBSD grant deployment in the area. The white markers represent the CBSDs, i.e., markers with dots are Category B and markers without dots are Category A. Note that the CBSD grant locations shown in the figures are only for illustration purposes due to proprietary concerns.

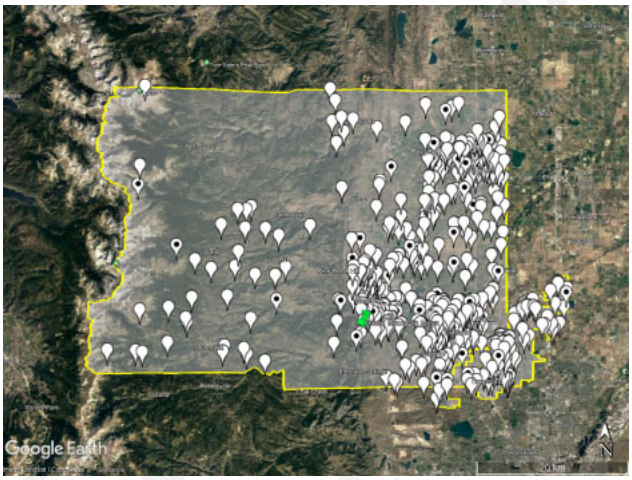

Figure 3.1: An example of a CBSD deployment in Boulder County and Broomfield County, Colorado.

The GMM sensor location is shown as a green pin in Figure 3.1. And Table 3.2 lists the parameters and values of the sensor. A 2D panel antenna is installed at 4.09 m height above ground level (AGL), with 12 degrees azimuth relative to True North and 0 degree down-tilt.

Green Mountain Mesa (GMM) Sensor	
Parameter	**Value**
Latitude	39.9918
Longitude	-105.274517
Height AGL (m)	4.09
2D Panel Antenna	Sector-KP-3SX4-90
Antenna Azimuth (degree)	12
Antenna Down-tilt (degree)	0

Figure 3.2: Sensor placement information at GMM, Boulder, Colorado.

3.1.2.2 Measured Power vs. Estimated Aggregated Emissions Comparison

Since there is no DPA near the Boulder area, an aggregate emission from all CBSD grants in Boulder and Broomfield Counties to the GMM sensor location is estimated for every 10 MHz channel in the 3550-3700 MHz band. An analogy is drawn between the estimated aggregate emission and the measured power by the prototype SEA sensor at the GMM location.

Figure 3.3 shows an example of median root mean square (RMS) power, max peak power, median RMS periodic frame power (PFP), and max peak PFP in 10 MHz channel collected by a NASCTN SEA prototype revision 3 sensor at the GMM location on December 6, 2022. More information on sensor data products can be found in Chapter 6. As shown in the figure, CBSD activities can be found in many channels (e.g., channels 1-4 and 11-14). Measurements in the last channel 15, centered at 3695 MHz, were in a sensor overload condition, which might affect the absolute power levels in that channel as well the 2 lower adjacent channels 13 and 14 which are centered at 3675 MHz and 3685 MHz, respectively.

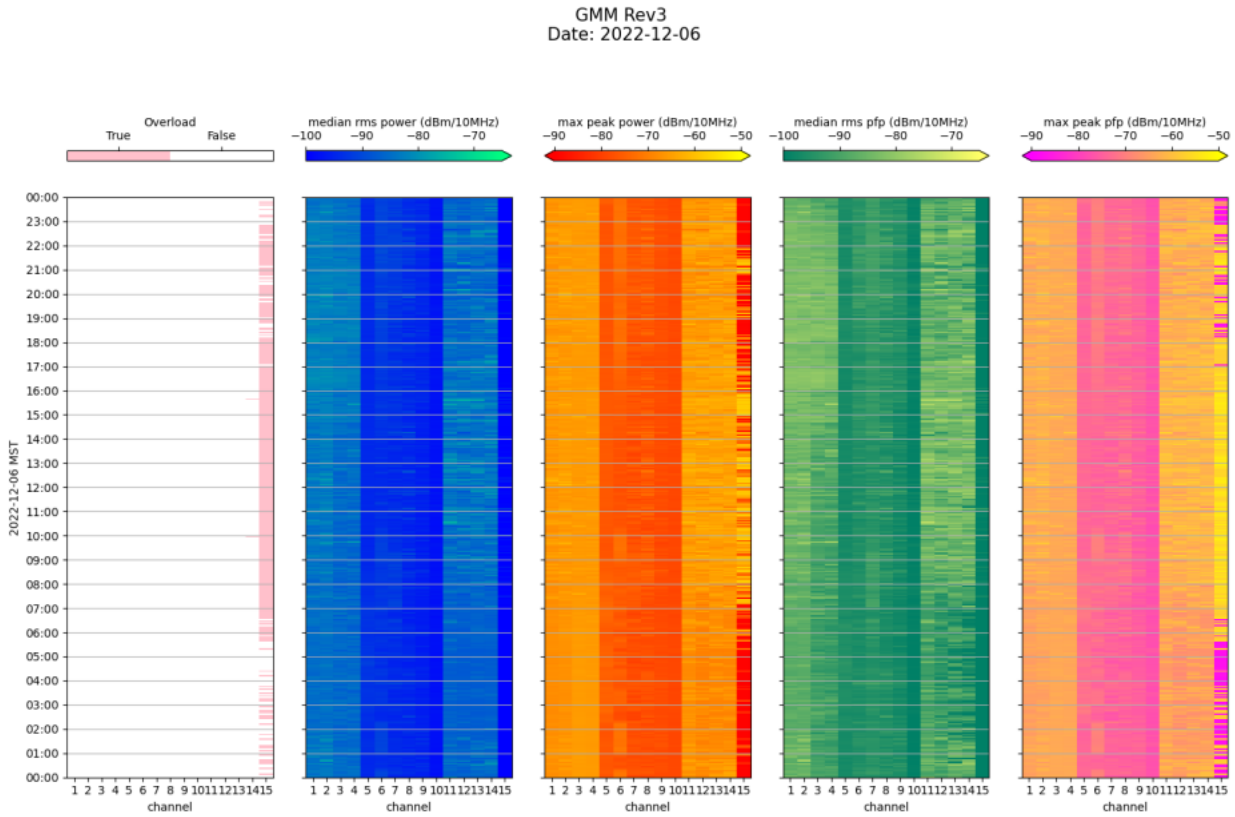

Figure 3.3: An example of measured power collected by a NASCTN prototype SEA sensor at the GMM location on December 6, 2022.

In addition, Figure 3.4 shows the median RMS power and max peak power collected by a NASCTN prototype SEA sensor at the GMM location from November 29, 2022 through December 7, 2022. The measurements were taken continuously with an on channel dwell time of 4 and channel revisit times of about 90 s, where each channel is measured once in a consecutive sweep.

On the other hand, Figure 3.5 illustrates an example of the estimated aggregate emission in 10 MHz channel from a randomized CBSD deployment to the GMM sensor location. We use different propagation models and configurations

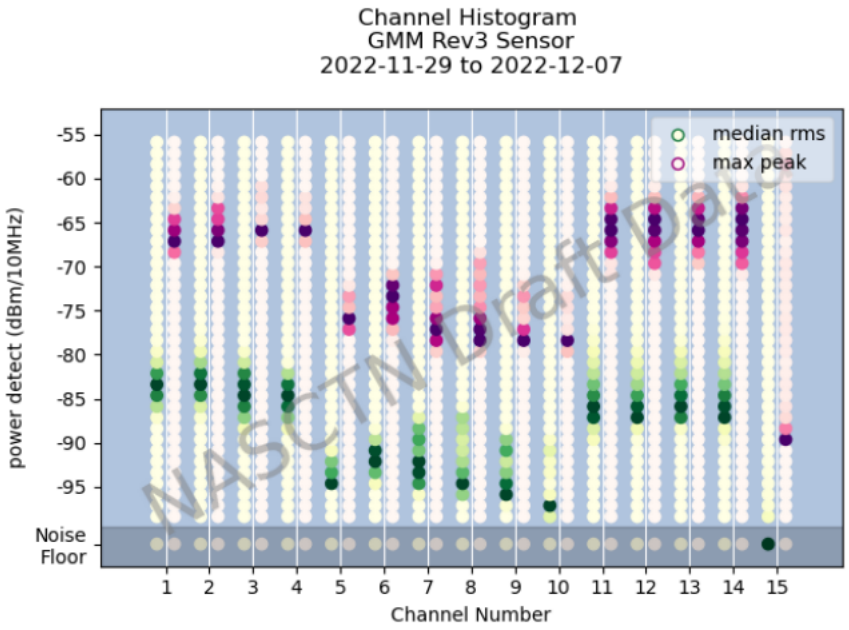

Figure 3.4: Measured powers collected by a NASCTN prototype SEA sensor at the GMM location from November 29, 2022 to December 7, 2022.

to estimate the path loss. In addition to the Irregular Terrain Model (ITM) model with 95th percentile from 2000 Monte Carlo iterations, the ITM model with median or reliability of 0.5 and the Friis free space model are used to compute the path loss from the CBSD to the GMM sensor location.

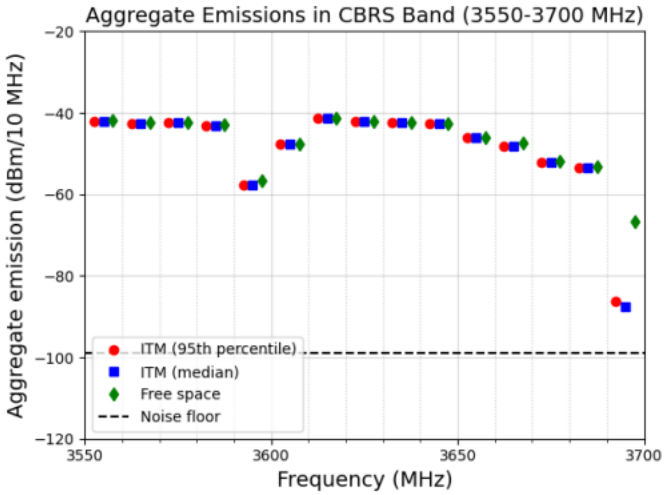

Figure 3.5: An example of estimated aggregate emission at the GMM sensor location.

The ITM model, with 95th percentile (as shown in red markers in the figure) and median (in blue markers), and the free space model (in green markers) provide similar aggregate emission results, except the last channel centered at 3695 MHz. To better visualize the emission results, the markers are shifted slightly in the x-axis but are still within the measured 10 MHz channel.

It is observed that the estimated aggregate emissions are much higher than the measured powers in all channels. It is found that the emissions from a few CBSDs, which are deployed within 3 km from the GMM sensor and in line-of-sight paths, dominate the aggregate emission levels. There are numerous factors that contribute to the discrepancy between the measured powers vs. estimated aggregate emissions such as

1. Inaccurate representation of the real-time CBSD deployment (e.g., location, antenna characteristics, transmit power, channel occupancy, etc.) generated from the quarterly obfuscated FAD data, and
2. Clutter has not been included in the aforementioned propagation models.

The aim is to determine these factors and resolve them so that the modeling results can be as close as possible to the real measurements. Getting more accurate representation of the real-time CBSD deployment could be feasible; however, adjudication and recommendations on propagation models and their feature set such as incorporating clutter is beyond the scope of the NASCTN CBRS SEA project.

3.1.2.3 Monte Carlo Simulation of CBSD Deployments

Monte Carlo simulation is used to determine if random CBSD deployments derived from the same obfuscated FAD data produce significantly different results. The obfuscated FAD data are obfuscated and aggregated from the FAD files by the SAS Administrators every quarter. The obfuscated FAD data have proven to be rich and informative, yet the data do not reflect real-time activities of CBSD deployment.

Figure 3.6 shows the estimated aggregate emissions for 100 different CBSD deployments. Since the census blocks near the GMM sensor are small in area, the CBSD locations do not vary substantially in effective path loss. It is also observed that the random deployments from the same obfuscated data produce similar aggregate emission results for both ITM and Friis free space models. This is likely due to the ITM model being a step wise model and heavily relies on the free space model for line-of-sight transmissions.

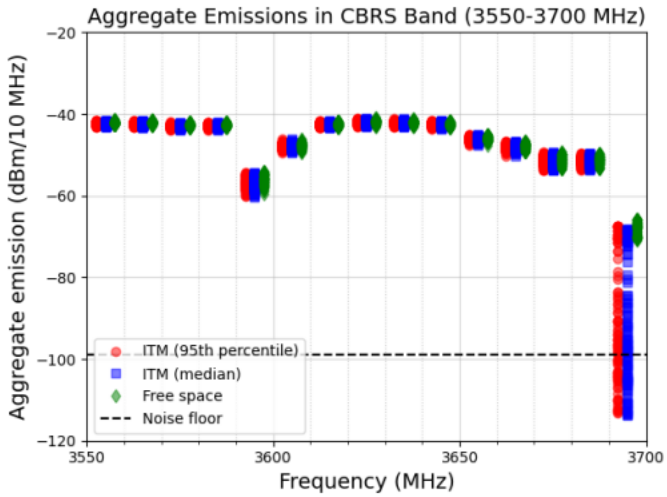

Figure 3.6: Examples of estimated aggregate emissions from 100 randomized CBSD deployments at the prototype GMM sensor in Boulder, Colorado.

The obfuscated FAD data could be sufficient for initial estimate of aggregate emissions at a given SEA sensor location. However, non-obfuscated FAD data with more detailed information can certainly improve the accuracy of aggregate emission calculation, which better informs site selections for SEA sensors.

3.2 SEA Sensor Siting

3.2.1 Site Criteria

The NASCTN SEA sensor locations are selected based on initial site surveys to ensure the availability of shelter and power as well as easy access for maintenance of the sensors over extended periods of time. High elevated sites are also preferable to better detect CBSD emissions at reduced obstructions and potential for clutter. Modeling is then utilized as a tool to further assist in selecting the sensor locations.

The noise level per 10 MHz channel at the sensor is computed as follows

$$Noise = -174\ (\text{dBm/Hz}) + 10\log_{10} BW + NF \tag{3.3}$$

The channel bandwidth BW is set to 10^7 Hz and the noise figure NF used for this modeling exercise is a projected 5 dB, which results in the noise floor to be equal to -99 dBm/10 MHz. Note that the noise figure used for modeling may differ slightly from the values of the deployed prototype systems as described in Section 4.4.

Three different 2D antenna patterns are used in the modeling as shown in Figure 3.7. They include a vertical polarized omnidirectional antenna, a slant polarized omnidirectional antenna, and a panel antenna. The omnidirectional antennas provide omnidirectional coverage in azimuth dimension and directive coverage in elevation dimension. Whereas, the panel antenna provides coverage from a planar array above the ground plane.

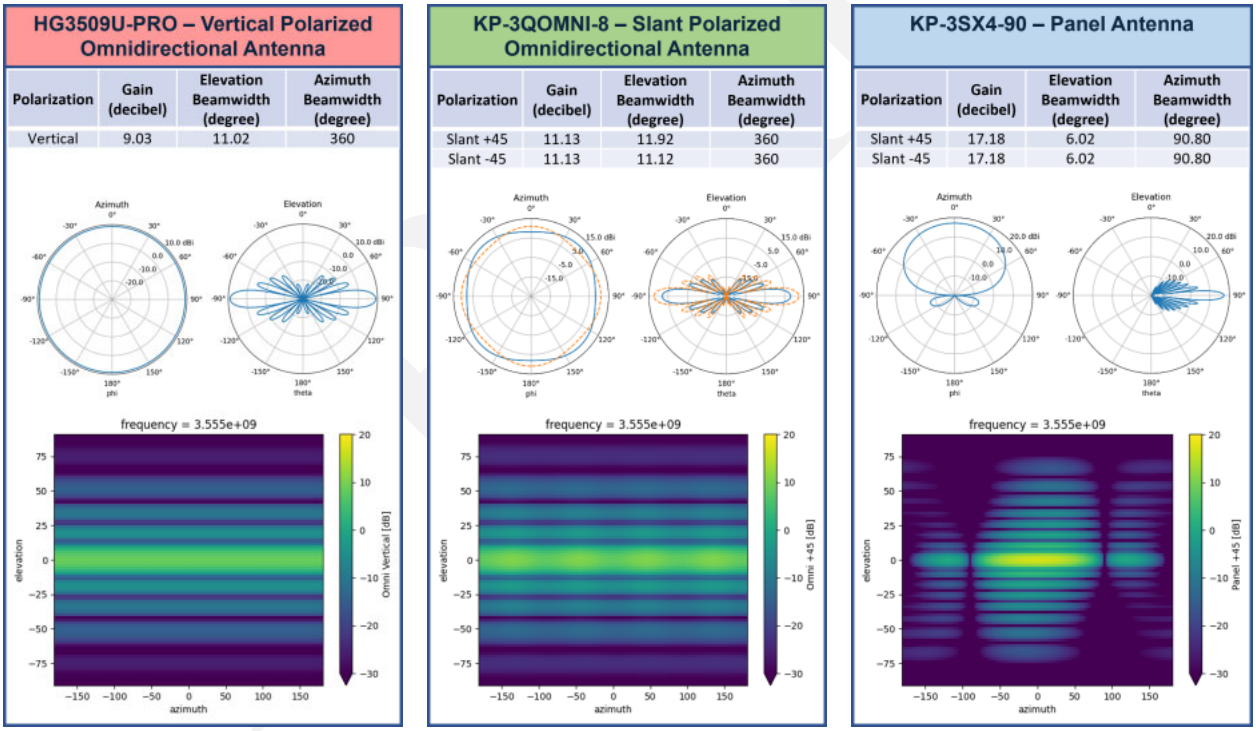

(a) Vertical polarized omnidirectional (b) Slant polarized omnidirectional (c) Panel antenna

Figure 3.7: Information of 2D antenna patterns used in the modeling.

3.2.2 Preliminary Aggregate Emission Results

The objective of the modeling in Task 2 is to estimate the aggregate emissions at candidate sensor locations inside the always-on DPA Camp Pendleton. The DPA is used to protect ground based incumbent systems operating below

3500 MHz. And the protection criteria, e.g., protection thresholds and neighborhood distances, vary by 10 MHz channels in the 3550-3650 MHz band [2]. Figure 3.8 depicts candidate sensor locations inside DPA Camp Pendleton. Most of these locations were recently surveyed by the NASCTN SEA team which will be described in detail in Section 8.2. The antenna height at each location is assumed to be 4 m AGL.

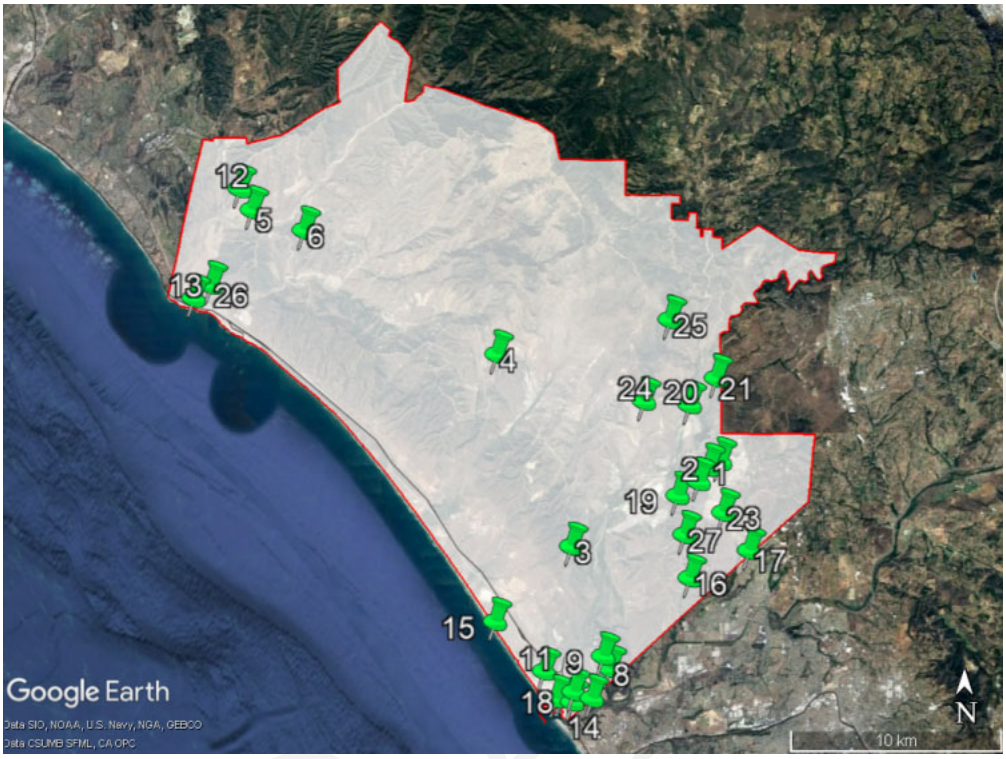

Figure 3.8: Candidate sites under investigation for simulation and field evaluation.

Figure 3.9 shows examples of the aggregate emission results at candidate sensor locations for the 3630-3640 MHz channel. Since Camp Pendleton is an always-on DPA, only emissions from the keep list are estimated. Although three 2D antenna patterns, described in Section 3.2.1, were used in the modeling, only the results for the vertical polarized omnidirectional antenna were presented as an example. The horizontal dashed line is the projected noise floor at -99 dBm/10 MHz. The emission levels are at or below the sensor internal noise floor at most candidate locations. In this example, for the 3630-3640 MHz channel, emissions are projected to be detectable above the noise floor only at candidate locations 8 and 27.

3.2.3 Key Technical Challenges with Current Simulation and Modeling

Several key technical challenges with the current simulation and modeling have been identified that hinder the accuracy of the modeling results. The preliminary modeling results presented in previous sections show that the aggregate emissions could be overestimated, as compared to the real measurements, potentially due to the following reasons:

- Maximum requested EIRP values of the CBSD grants have been used in the modeling, while in practice, approved CBSD grants may operate at lower maximum allowable EIRP values or may not transmit at all.
- The effects of activation of adjacent DPAs or other protection entities nearby, e.g., Environmental Sensing Capability (ESC) protection zones, have not been considered in the modeling.
- Clutter has not been taken into account in the propagation models used to compute path loss from the CBSD grant to the sensor location.

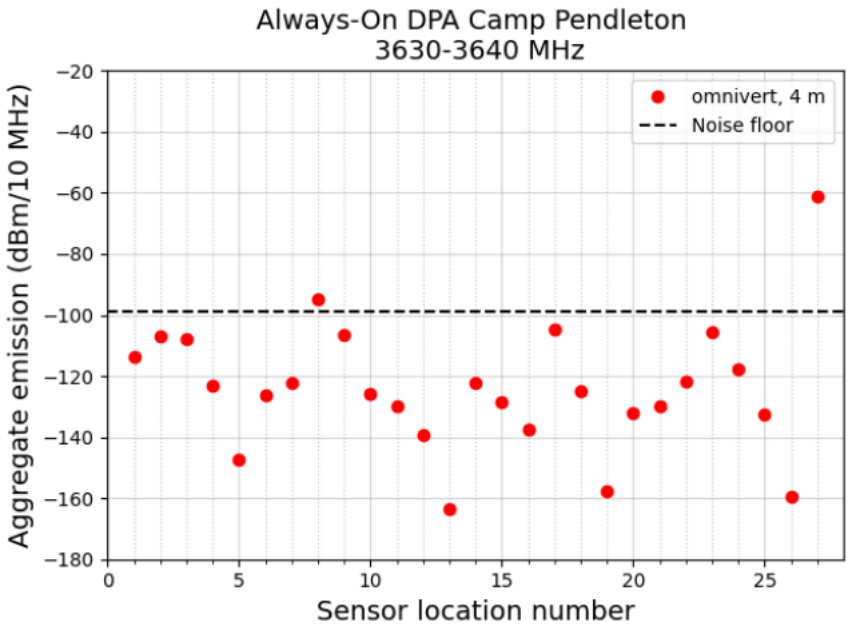

Figure 3.9: Simulated aggregate emission results at candidate sensor locations inside DPA Camp Pendleton.

3.3 Informing Data Products

This section sketches out some initial ideas to aid with the data products development. Chapter 6 describes the other data products in detail. Data products and data analysis are used to a) present the aggregate emissions in the CBRS band due to deployment of systems over time, and b) validate the always-on DPA protection area at Camp Pendleton.

For the DPA Camp Pendleton, the protection threshold $threshold_{ch}$ for each 10 MHz channel ch can be obtained from the GB-DPAS.kml file at [2]. Let's define $P_{s,t,ch}$ (dBm/10 MHz) as the aggregate emission measured by sensor s at time t in channel ch. Figure 3.10 illustrates an example of the measured emissions over time, as compared with the threshold $threshold_{ch}$, by different sensors inside the DPA. The time resolution Δt_s at sensor s can be varied depending on the analysis types and capabilities of the sensor.

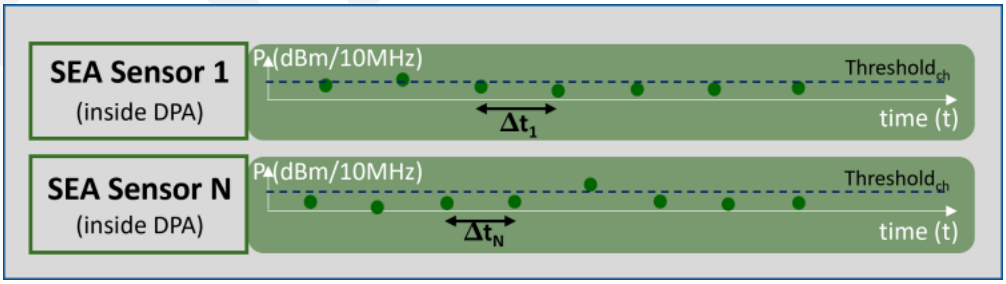

Figure 3.10: Example of measured emissions over time by different sensors inside the always-on DPA Camp Pendleton.

An alternative approach is to show the histograms of mean measured emission per a period, e.g., hour or day, over a longer period of time, e.g., year, for all channels as shown in Figure3.11.

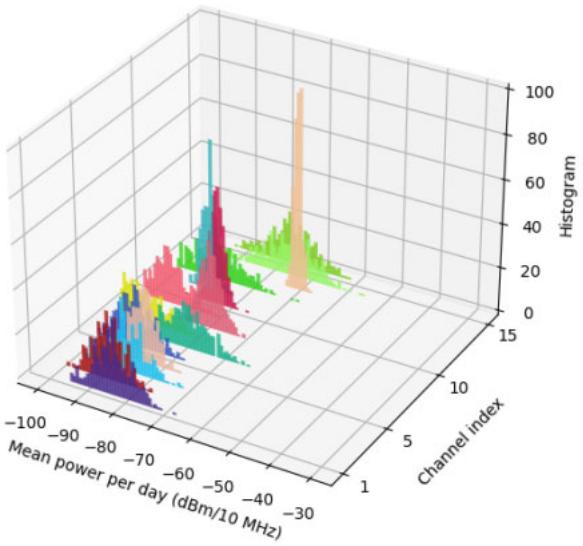

Figure 3.11: Example of statistics of mean power per day over year for all channels at a single sensor.

To assess the measured aggregate emission $P_{s,t,ch}$ against the protection threshold $threshold_{ch}$ on channel ch, a simple algorithm as shown in Figure 3.12 can be carried out. Note that, conversion of the protection threshold $threshold_{ch}$ might be required, since the threshold obtained from the GB-DPAS.kml file was derived by the National Telecommunications and Information Administration (NTIA) with different receive system assumptions.

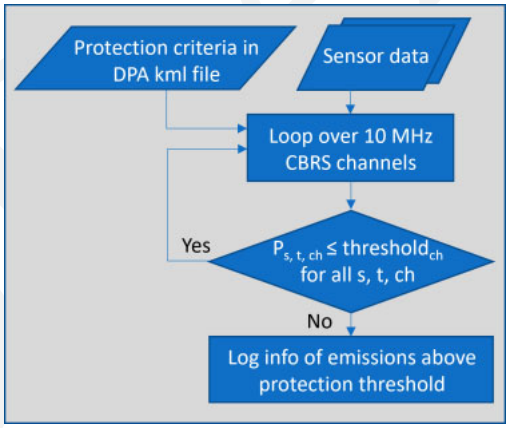

Figure 3.12: Example of a workflow used to assess aggregate emissions in the CBRS band.

Chapter 4

Sensor System

4.1 Sensor Spiral Development Process

A custom sensor has been developed to monitor and analyze radio frequency (RF) emissions in the Citizens Broadband Radio Service (CBRS) band. The data produced by a network of these sensors supports Sharing Ecosystem Assessment (SEA) tasks 1 and 2. A spiral development process was utilized to rapidly provide benchtop prototypes, allow incremental development as system requirements solidified, and support the evolution of the spectrum-characterization and occupancy sensing (SCOS) framework. The development process is comprised of 5 stages as shown in Figure 4.1, and the first three prototype phases will inform the final two design iterations.

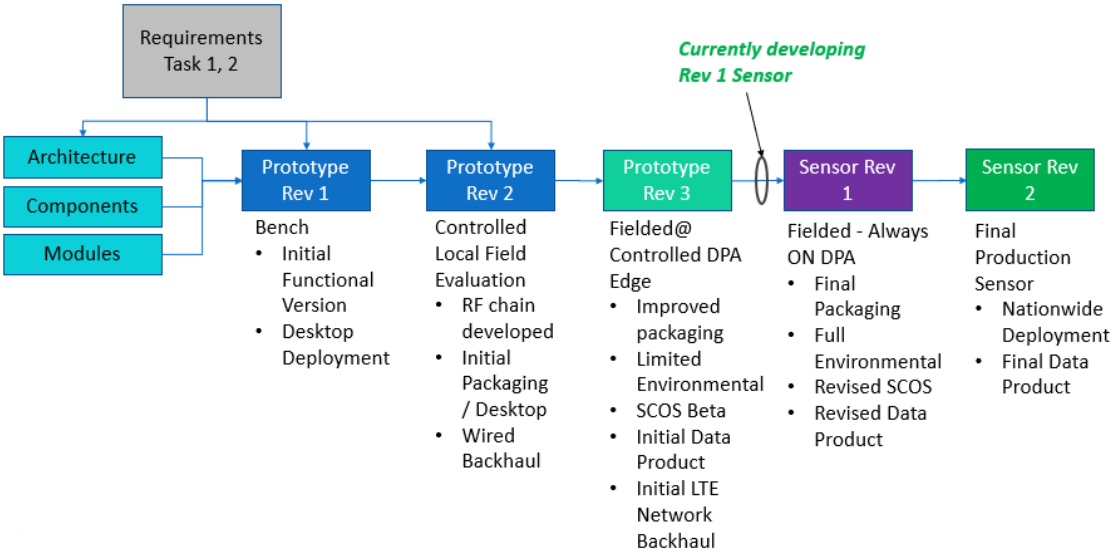

Figure 4.1: Spiral Development Process

4.2 Sensor Overview

The SEA sensor was designed to analyze RF emissions in the 3550 to 3700 MHz CBRS band. It is a complete system that captures the RF spectrum, computes power statistics and their distributions across time and frequency, and transfers these data products to a centralized repository for post-analysis. The following sections provide further details on the sensor design, and refer specifically to the final prototype revision of the system (revision 3) unless otherwise noted.

4.2.1 RF System

Figure 4.2 is a block diagram of the complete sensor RF system. At the start of the RF signal path is a single pole double throw (SPDT) switch (designator SWT in block diagram) which selects either the antenna (ANT) or a calibrated noise source (NSR) as the input to the system. The antenna receives RF emissions in the CBRS band, while the noise source is used to perform periodic calibration as described in Section 4.2.3. Following the RF switch are an isolator (ISL) and a bandpass filter (BPF). The filter rejects emissions outside of the CBRS band, while the isolator helps to present a consistent impedance to the input of the filter regardless of the state of the RF switch. This minimizes a variable between the calibration path and antenna path, and ensures a more accurate calibration. Following the filter is an RF limiter (LMT) to provide protection against strong in-band emissions. Next, a low noise amplifier (LNA) (AMP1) provides the gain needed to compensate for system losses and minimize system noise figure.

Following the LNA are a coaxial cable (CBL) and attenuator (ATT). A coaxial cable up to 30.5 m long is accounted for in the system gain budget. This allows for physical separation between groups of components, which provides

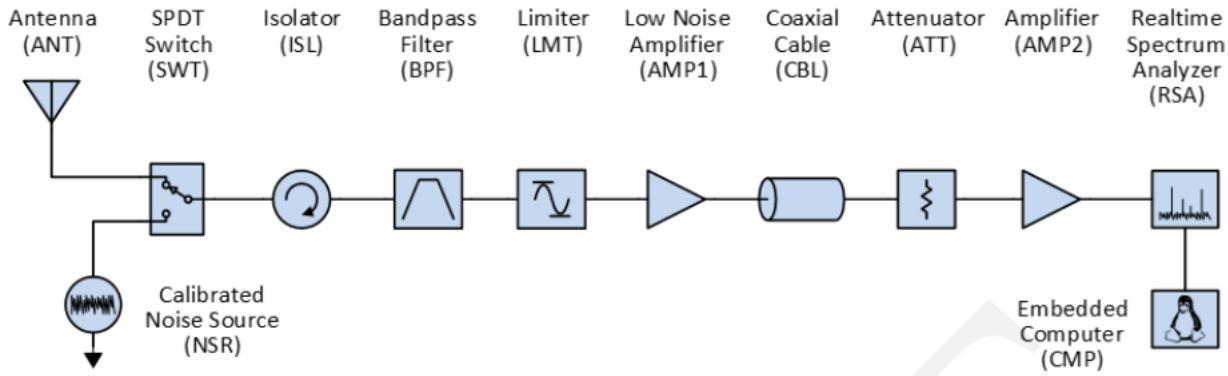

Figure 4.2: Sensor RF System Diagram

flexibility in how the sensor is deployed. The group of components prior to and including the LNA (AMP1) can be close to the antenna to minimize system noise figure, while the remaining components which tend to be larger and consume more power can be located farther away. The attenuator following the coaxial cable is used to make adjustments to compensate for cables of varying lengths. The attenuator value is adjusted as needed to maintain a constant net insertion loss, resulting in consistent RF performance regardless of system variations due to installation.

The final components in the RF signal path are the second amplifier stage (AMP2) and real-time spectrum analyzer (RSA). The second amplifier stage preceding the RSA provides additional gain needed maintain a low overall system noise figure. While the RSA can operate as a fully-featured and stand-alone test instrument, in this application it serves as a tunable RF digitizer. The RF signal at the input is converted to a digital complex baseband signal with in-phase and quadrature (IQ) components. Subsequent processing on the IQ data is performed on the embedded computer (CMP).

Tables 4.2 and 4.3 provide a complete list of components and part numbers used in the sensor. Additional details on select components are provided in the following sections.

4.2.1.1 Antenna (ANT)

The sensor uses an L-Comm HG3509U-PRO [10] antenna. This antenna is vertically polarized with 9 dBi gain and has a specified operational bandwidth of 3400 to 3700 MHz. The radiation pattern is omnidirectional in azimuth and has an 11° 3 dB beamwidth in elevation. This pattern is advantageous because an omnidirectional pattern allows the sensor to observe in a 360° field of regard in azimuth. Since CBRS systems are terrestrial systems, a more directive elevation pattern is useful as neighboring Citizens Broadband Radio Service devices (CBSDs) will be predominately located within the 3 dB beamwidth of the sensor antenna's elevation pattern.

The sensor has a single input for the RF path. For simplicity, a vertically polarized antenna is leveraged to observe RF emissions in the CBRS band. CBSDs typically leverage dual slant 45° polarization. The measurement system may incur up to a 3 dB polarization loss from the measured dual slant 45° and circularly polarized emitters in the CBRS ecosystem. An additional polarization loss will be incurred from horizontally polarized emitters broadcasting in the measurement band.

4.2.1.2 Bandpass Filter (BPF)

A custom-designed cavity filter from K&L Microwave, part number 13FV40-3625/U150-O/O, is used to protect the LNA from potentially strong out-of-band emissions while allowing the CBRS band to pass with minimal insertion loss. Since the filter is located prior to the LNA (AMP1), it's insertion loss directly impacts the overall system noise figure. The filter was specified to have an insertion loss of less than 1.5 dB at the center frequency and no more than

Table 4.1: Measured Filter Response Characteristics

Insertion Loss (3625 MHz)	-1.3 dB
CBRS Band Flatness (3550-3700 MHz)	0.85 dB

Filter Shape			
Parameter	**Low**	**High**	**Bandwidth**
1 dB Bandwidth	3545 MHz	3705 MHz	160 MHz
3 dB Bandwidth	3535 MHz	3718 MHz	183 MHz
>20 dB Rejection	<3528 MHz	>3726 MHz	198 MHz
>40 dB Rejection	<3519 MHz	>3734 MHz	215 MHz
>60 dB Rejection	<3505 MHz	>3747 MHz	242 MHz

1 dB roll-off across the range of CBRS frequencies, while providing suppression of at least 57 dB for all frequencies either 50 MHz below or 50 MHz above the CBRS band.

The realized filter was characterized using a Rohde & Schwarz ZNB8 vector network analyzer and the resulting insertion loss, shown in Figure 4.3, indicates that it meets these specifications. A summary of measured filter response characteristics is provided in Table 4.1.

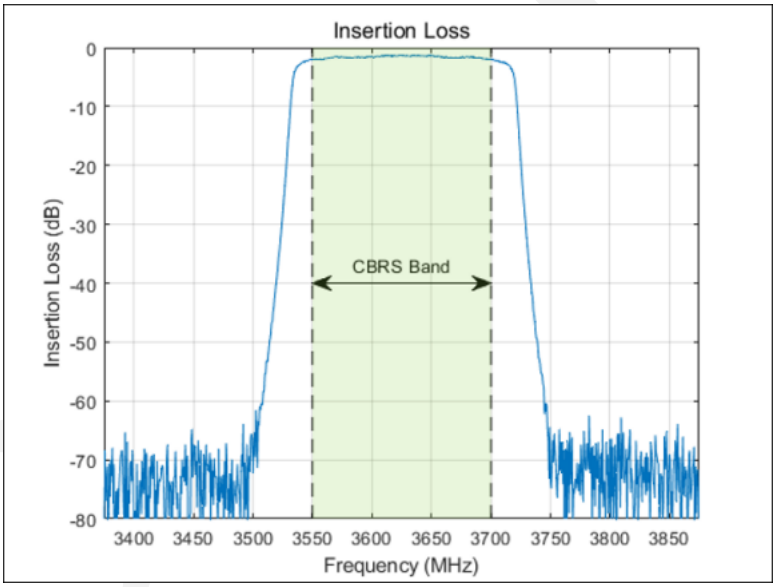

Figure 4.3: Filter Measured Insertion Loss

4.2.1.3 Low Noise Amplifier (AMP1)

The first LNA in the receive chain of the measurement sensor is the LNA02004000A, 0.2 GHz to 4.0 GHz wideband low noise amplifier from LiConn, Inc. This LNA has a typical gain of 29 dB as shown in Figure 4.4, and a typical noise figure of 1.3 dB as shown in Figure 4.5 [1]. The high gain allows for the noise figure to be maintained between the portion of the RF sensor co-located with the antenna and the input of the RSA. This allows additional flexibility for system deployment. The RSA may be located over 30.5 m away from the antenna without degradation to RF performance of the system. This LNA is rated to have a wide operational temperature of -45°C to +85°C which simplifies the enclosure design by not requiring active heating or cooling.

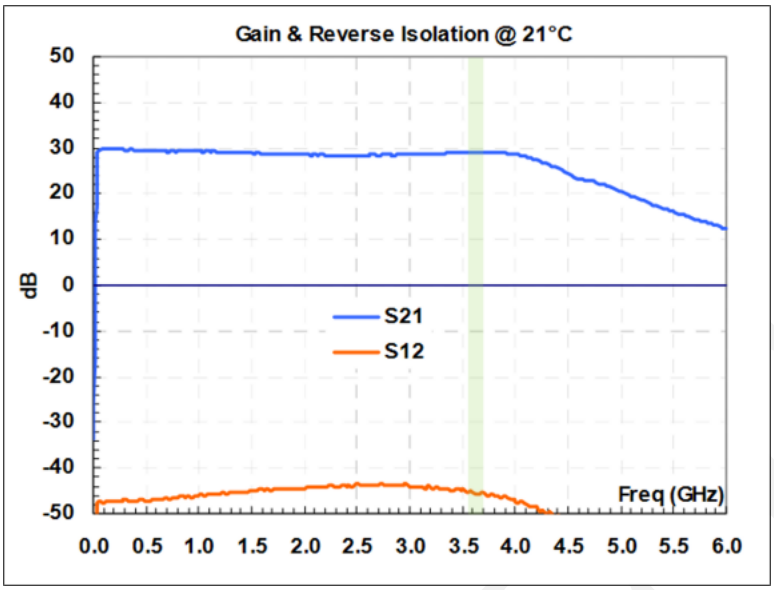

Figure 4.4: LNA Gain and Reverse Isolation from Manufacturer Datasheet [1]

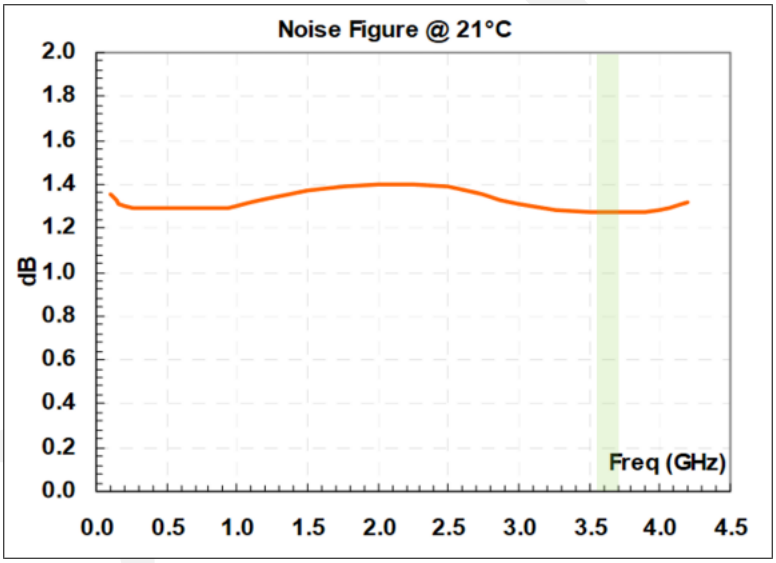

Figure 4.5: LNA Noise Figure from Manufacturer Datasheet [1]

4.2.1.4 Realtime Spectrum Analyzer (RSA)

A market survey was conducted to identify RSAs, spectrum monitoring receivers, and software-defined radios that are tunable over the CBRS band, can acquire at least 20 MHz instantaneous bandwidth and can transfer baseband IQ data over a universal serial bus (USB) or Ethernet interface. Thirteen candidates were identified and evaluated based on RFs performance, size, power consumption, operating temperature range and cost. Of those thirteen units, three units were selected to undergo laboratory testing to examine RF performance, and ultimately the Tektronix RSA507A was selected.

The RSA507A is tunable over the frequency range 9 kHz to 7.5 GHz and has an instantaneous bandwidth of up to

40 MHz. It has a USB 3.0 interface for device configuration and for transferring IQ data to the host computer when operating in IQ acquisition mode. It draws 15 W, operates over the temperature range -10°C to +55°C and comes in a 30.0 cm x 27.2 cm x 6.9 cm rugged package.

Laboratory measurements were performed to verify critical datasheet specifications [11] and to examine RF performance under the expected operating conditions for the SEA sensor application. Several of these measurements are presented here.

First, the amplitude variation and spur levels over the RSA's maximum instantaneous bandwidth of 40 MHz were measured. The instrument was configured with a center frequency of 3600 MHz, reference level of -40 dBm, and IQ acquisition sample rate of 56 Msps (which, by design, is the sample rate corresponding to 40 MHz usable bandwidth). An RF signal generator was then used to sweep a -40 dBm continuous-wave (CW) signal across the band in 1 MHz steps, and IQ data was acquired at each 1 MHz step. A fast Fourier transform (FFT) analysis was performed on these IQ captures and the fundamental signal level and spur levels were measured. Figure 4.6 shows an example of the FFT analysis performed on the capture when the input frequency was 3605 MHz. The green marker indicates the frequency and power level of the input signal, and the red markers indicate the frequencies and power levels of the strongest three spurs. Note that with an IQ sample rate of 56 Msps, the baseband FFT has a span of -28 MHz to +28 MHz. However, only the middle 40 MHz, or a span of -20 MHz to +20 MHz, is intended to be usable bandwidth. The frequency axis in Figure 4.6 has been modified to indicate the RF input frequency rather than the baseband frequency of the acquisition, and frequencies outside of the usable bandwidth have been shaded orange.

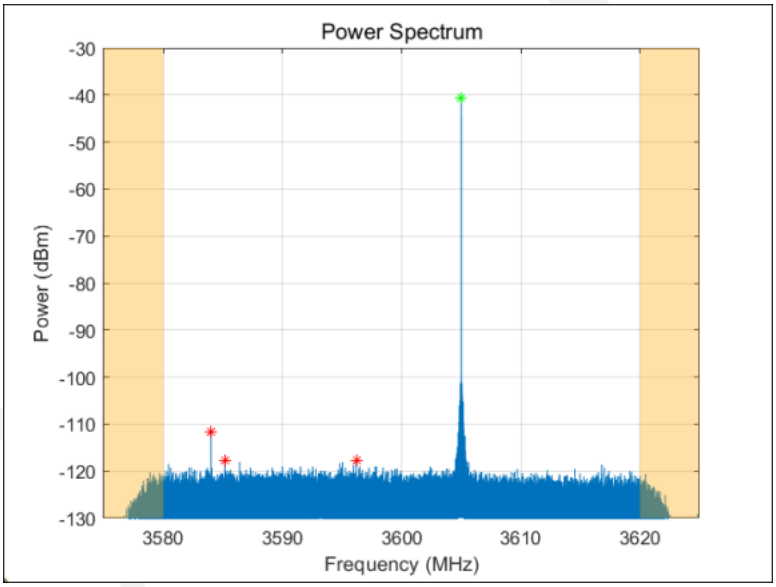

Figure 4.6: RSA Power Spectrum with 3605 MHz CW Input Signal

Figure 4.7 shows the resulting fundamental signal level and relative spur levels as the input signal is swept across the band. The plot of the fundamental signal level illustrates an amplitude variation of approximately 0.2 dB across the 40 MHz band. Note that conclusions about amplitude accuracy should not be drawn from this plot because the insertion loss of the short coaxial cable between the signal generator and RSA was only estimated. The plot of the spur levels shows that for all input frequencies, the largest spur remains at least 70 dB below the fundamental. The result is shown in Figure 4.8.

Next, to characterize the noise performance of the RSA, its input was terminated with a 50Ω load and IQ captures were obtained. With the RSA still configured for center frequency of 3600 MHz and sample rate of 56 Msps, and

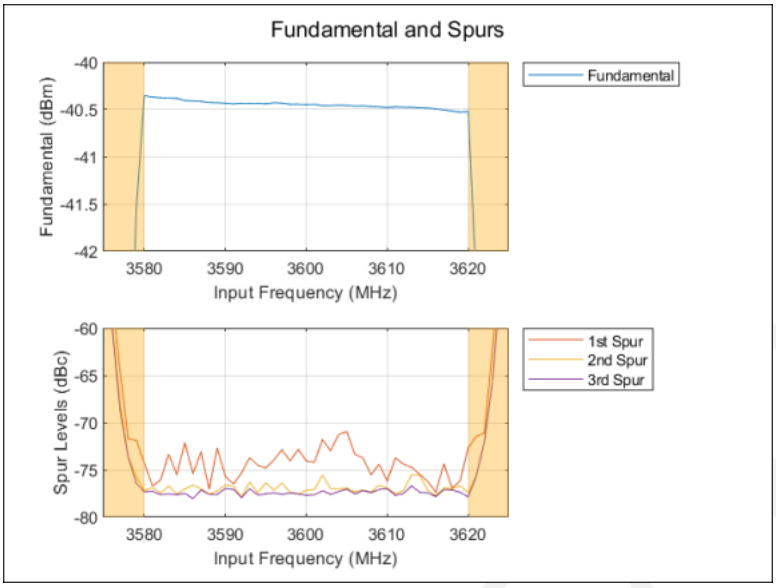

Figure 4.7: RSA Fundamental Signal and Spur Levels

with the internal preamplifier enabled, the RSA reference level was increased from -54 dBm to 0 dBm in 6 dB steps. At all reference level settings, the RSA's internal input attenuator was set to 0 dB to optimize noise performance. IQ data was acquired at each reference level setting, and from the IQ data the displayed average noise level (DANL) was computed. The result is shown in Figure 4.8. The minimum measured DANL of -155.6 dBm is within the range specified in the datasheet [11].

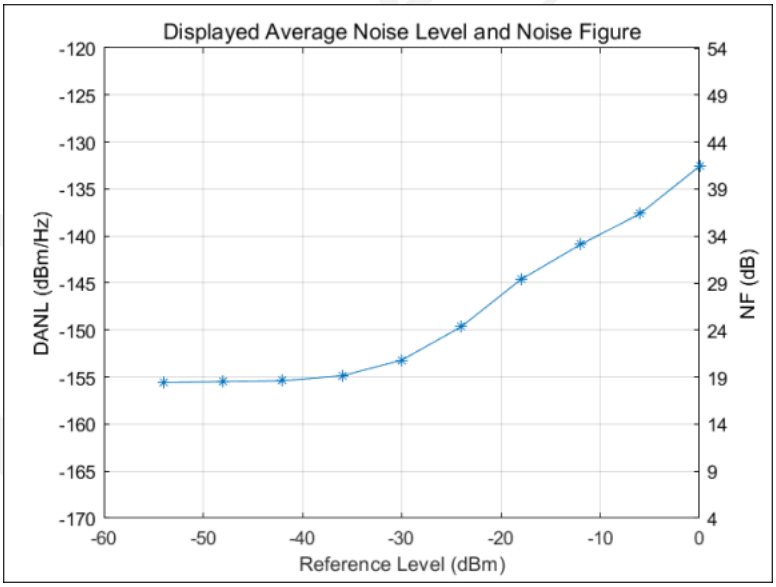

Figure 4.8: RSA Displayed Average Noise Level and Noise Figure

As noted on the right axis of Figure 4.8, the RSA's noise figure (NF) is equivalent to the difference between DANL and thermal noise power (nominally -174 dBm/Hz for the expected operating temperature range). It is important to know the RSA's noise figure for its given reference level to determine the overall system noise figure.

The gain compression of the RSA was also measured. With the RSA still configured for center frequency of 3600 MHz and sample rate of 56 Msps, and with the internal preamplifier enabled, an RF signal generator was used to produce a 3600 MHz CW signal at increasing power levels. At each power level, IQ data was acquired and FFT analysis was performed to determine the signal level measured by the RSA. The 1 dB gain compression (P1dB) point was determined by observing the measured signal level over the range of known input levels.

An example P1dB measurement made with an RSA reference level of -30 dBm is shown in Figure 4.9. In this plot the red points indicate input signal levels where the RSA reported an "ADC overrange" condition. For all points where analog-to-digital converter (ADC) overrange does not occur, the measured signal power closely matches the input power and no gain compression is exhibited. When an ADC overrange condition does occur, large distortion products in the power spectrum draw power away from the fundamental, and as a result the measured fundamental signal level drops abruptly. Due to this characteristic, the P1dB point is defined to be the highest input power that does not cause an ADC overrange condition. While Figure 4.9 shows the gain compression measurement with an RSA reference level of -30 dBm, this characteristic is typical of most reference levels where this measurement was performed.

The process of measuring P1dB was repeated for a range of RSA reference level settings. The resulting P1dB over the range of RSA reference levels is shown in Figure 4.10.

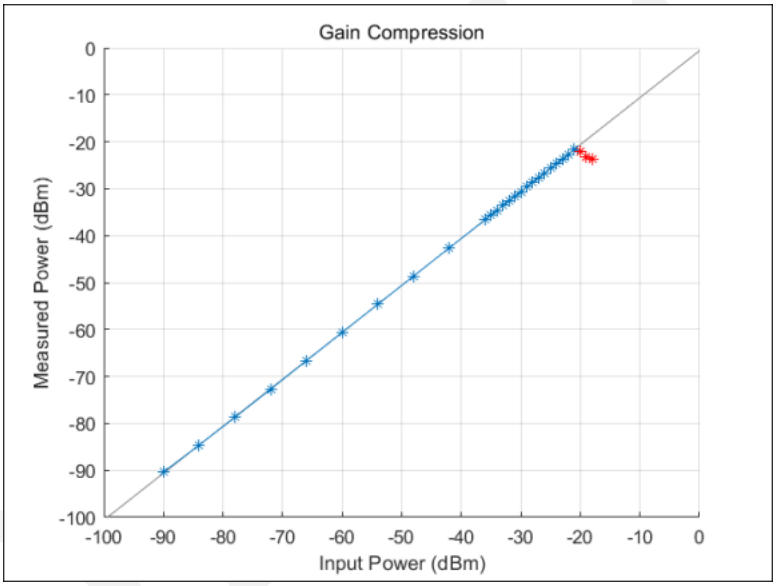

Figure 4.9: RSA Gain Compression Measurement Example

4.2.1.5 Embedded Computer (CMP)

Processing of the IQ data is performed on a SimplyNUC NUC11TNHI7Z small form factor computer. This computer has an 11th Gen Intel® Core™ i7 processor and 32 GB RAM, and runs the Ubuntu operating system. The computer ingests IQ data from the RSA via a USB 3.0 interface and performs processing as described in Section 6.1.2.

4.2.2 Cascade RF Analysis

System cascade analysis has been generated to optimize the RF chain to the signal characteristics expected at the antenna input of the sensor. The goals are to optimize the signal levels (min/max) in the signal processing chain while providing a low overall noise figure by adding a minimal amount of noise to the signal. Once the RF signal is digitized in the signal processor, the overall signal profile is frozen, so it is critical to optimize the signal before it is passed on to the digitizer in the signal analyzer. To design the system, components need to be selected to address overall system constraints and goals including:

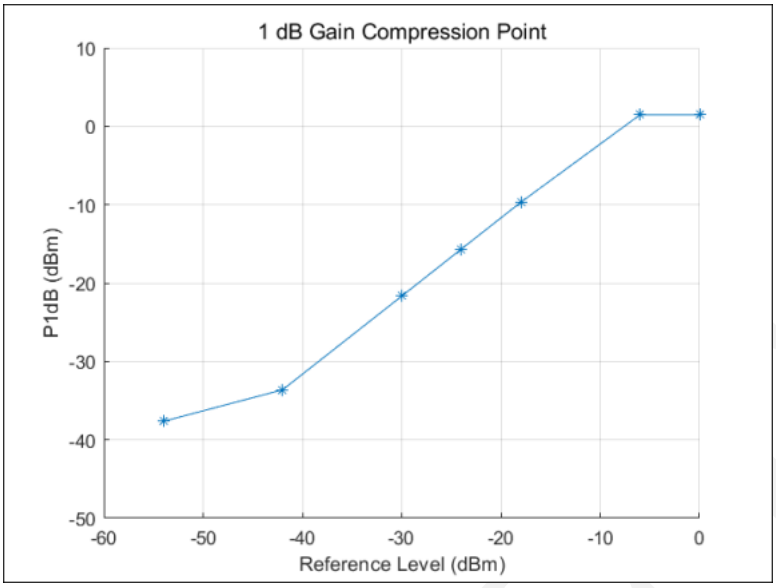

Figure 4.10: RSA 1 dB Gain Compression Point

- Minimum noise figure for sensitive measurements
- Output power level needs to be controlled to not compress the signal analyzer
- Need to maximize dynamic range
- Survive maximal signal levels
- Filter external signals outside of the CBRS band
- Variable distance between preselector and signal processing unit

The constraints are balanced by trading performance characteristics such as gain, compression (linearity), noise figure maintenance, and signal level. Initially, this design is accomplished by cascading system components utilizing information provided by component data sheets, and the predicted performance relies on this information. Cascaded analysis provides insight into internal signal characteristics with overall performance in mind.

In a system cascade, the noise figure is dominated by initial components that are in front of the first LNA gain stage. The design of the preselector includes a calibration switch and preselection filter. Additionally, a power limiter has been included to provide protection to the LNA in the event of a large signal being present at the input. The design of the signal processing unit with the added LNA is such that a fixed loss can be accommodated between the preselector and signal processing unit. This enables a long cable to be placed between the preselector and the signal processing unit for flexibility in installation locations. An added attenuator is included in the signal processing unit to ensure the LNA and signal analyzer do not compress the signal. A full system block diagram is shown in Fig. 4.11. The gain of full system is 32.1 dB and the noise figure is 4.9 dB as shown in Fig. 4.12. The dominant factor in this overall noise figure is due to loss from the components before the LNA stage.

Finally, utilizing the cascaded gain and noise figure for the system is shown versus frequency in Fig. 4.14 demonstrating performance across the entire CBRS band.

4.2.3 On-board Calibration

On-board calibration is integrated into the measurement system. A NIST-traceable noise diode (NSR) is integrated into the front of of the measurement system. The purpose of the calibration is to minimize the uncertainty of the measurement. This allows the system to compensate for system variations of components over time. An example

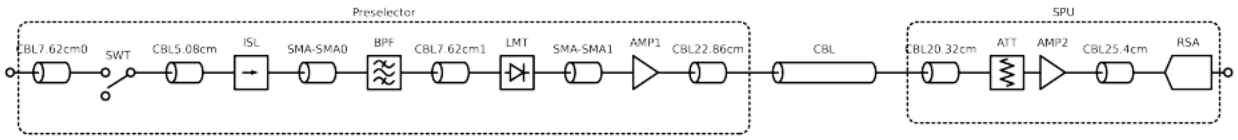

Figure 4.11: Full system cascade block diagram illustrating the components and interconnects that are considered in estimating the cascaded gain and noise figure.

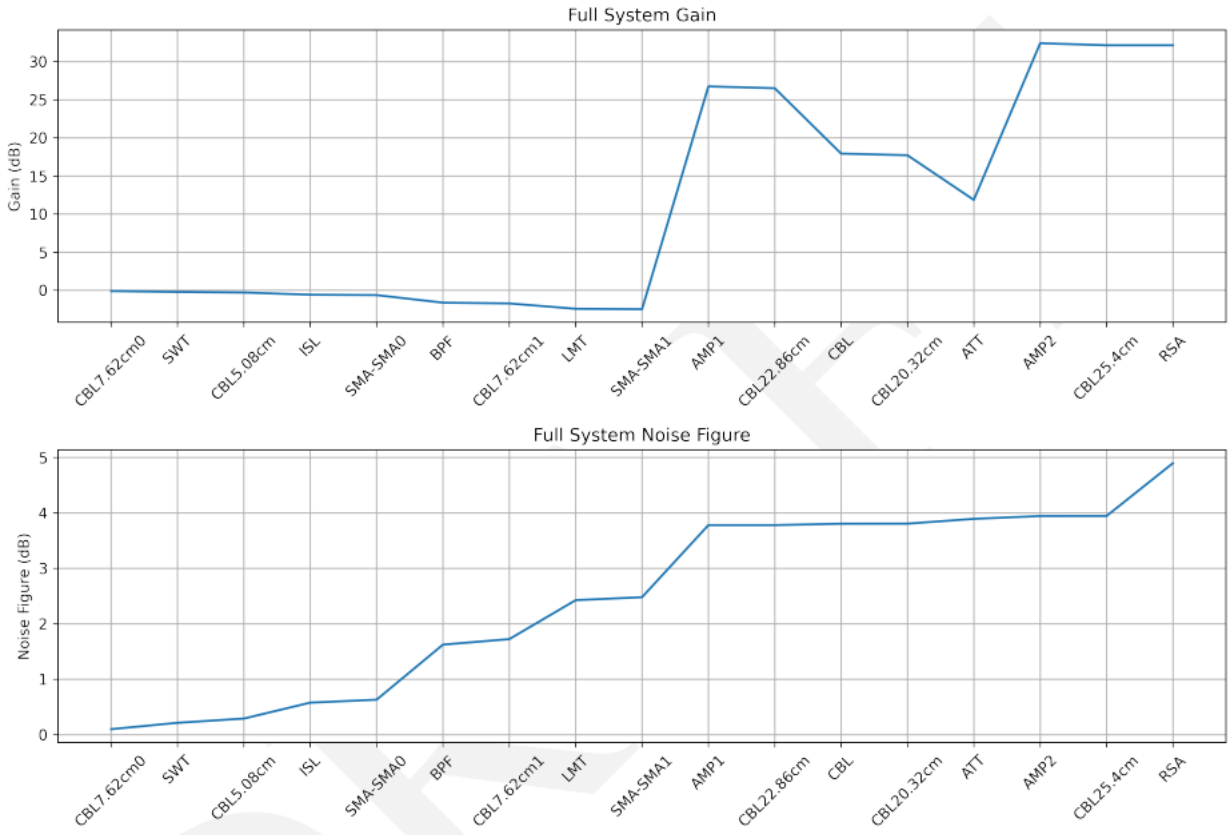

Figure 4.12: Cascaded gain (top) and noise figure (bottom) of the full system showing the contribution of each component to the overall system gain of 32.1 dB and noise figure of 4.9 dB.

of this effect is shown in the NTIA TR-20-548 [12]. In this report, the gain of the measurement sensor shown in Figure 4.15, referred to as the preselector, shows gain drift over a multi-year period. The gain of the system lowers as a function of time. These macro changes, as well as diurnal changes due to temperature variation, can be accounted for by leveraging an on-board calibration source.

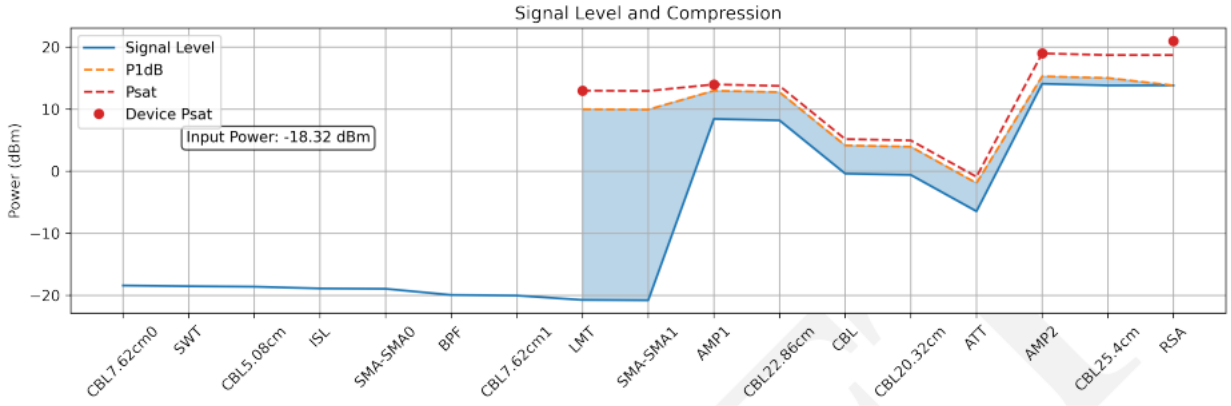

Figure 4.13: Cascaded signal level and compression characteristics of the full system when driven at the P1dB compression point of -18.32 dBm. This demonstrates how the signal compares to the compression point at each point in the system.

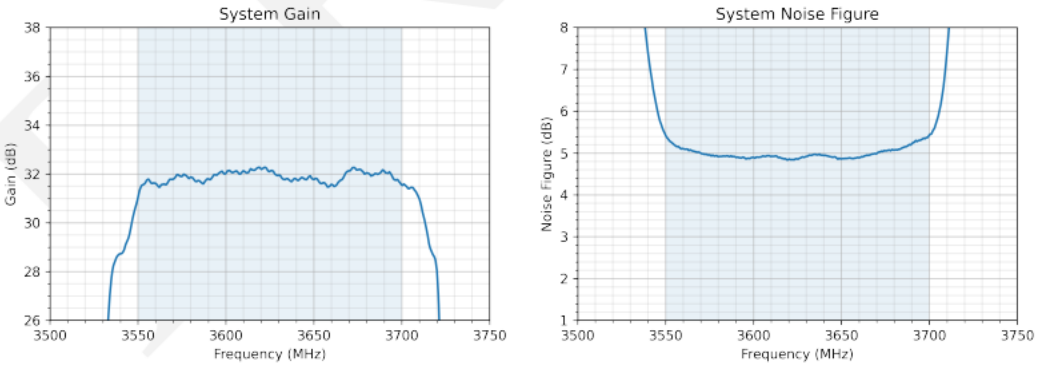

Figure 4.14: Cascaded gain and noise figure prediction of the system over the entire CBRS band.

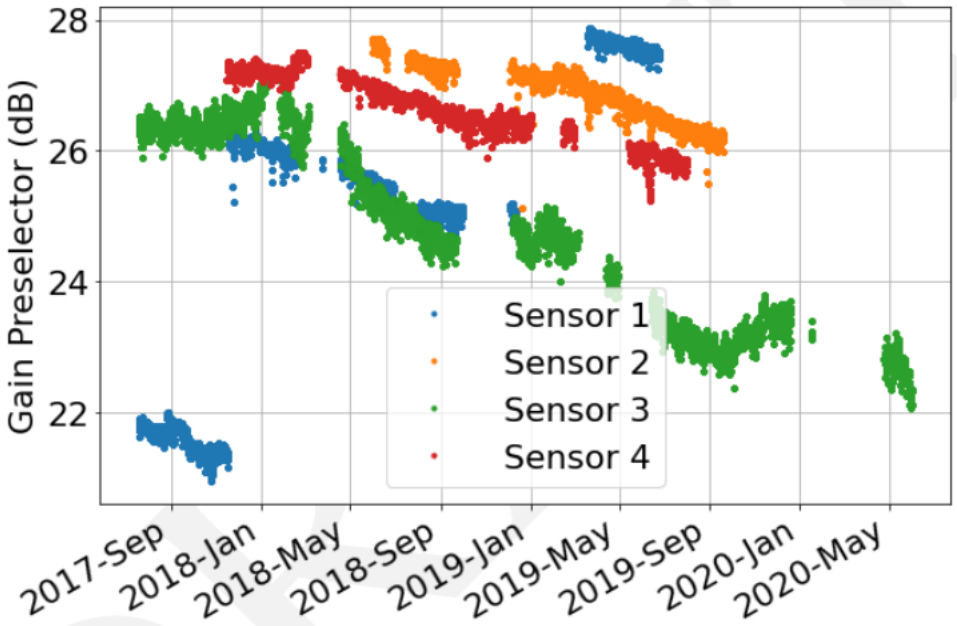

Figure 4.15: Gain drift shown of measurement sensor leveraged in NTIA's TR-20-548 Technical Report.

4.3 Sensor Implementation

Since certain RF components need to be located near the antenna to minimize system noise figure, while other high size, weight and power (SWaP) components can be placed farther away, the system is partitioned into two main assemblies to add flexibility to support various field installation requirements. These assemblies are the preselector and the signal processing unit (SPU). The preselector assembly includes the RF components up to and including the LNA (designator AMP1 in Figure 4.2), along with supporting control and power distribution components. These components dissipate relatively little power and can tolerate an operating temperature of at least 50°C. As such, the preselector is assembled into a fully sealed weatherproof enclosure that can be mounted on a mast directly below the antenna. The SPU assembly, on the other hand, contains components which dissipate more power and are generally larger and heavier. However, these components can be located farther from the antenna without sacrificing overall system performance. Therefore, the current version of the SPU is built with the intention of being installed in a climate-controlled area where size and weight are not of primary concern, and can interface with the preselector via cables with lengths up to 30.5 m.

Figure 4.16 is a high-level block diagram of the overall system as it has been implemented. In addition to the aforementioned preselector and the SPU assemblies, this diagram includes the commercial off-the-shelf (COTS) antenna and long-term evolution (LTE) router, which enables remote access to the system.

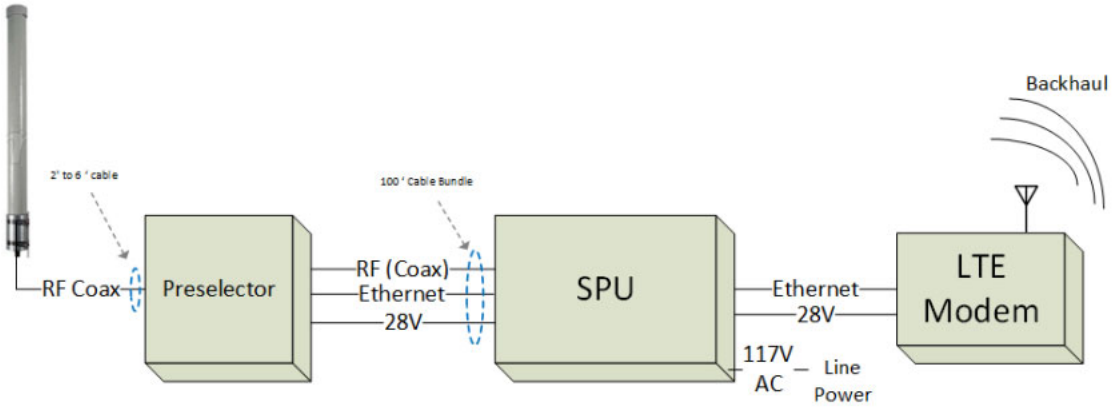

Figure 4.16: SEA Sensor Block Diagram

The following sections provide further details on the preselector and SPU assemblies.

4.3.1 Preselector Assembly

The preselector assembly contains all RF system components up to and including the LNA (AMP1), along with supporting power distribution, control and environmental monitoring components. A block diagram of the preselector is provided in Figure 4.17 and component part numbers are provided in Table 4.2. Note that Table 4.2 also includes component operating temperature ranges, which will be referenced in the following subsection.

A network-enabled relay module in the preselector (reference designator CTL1 in Figure 4.17) allows the RF switch (SWT) state, power to the LNA (AMP1), and power to the noise source (NSR) to be controlled remotely. Overall system power distribution is presented in more detail in Section 4.3.3. Additionally, the relay module supports the remote monitoring of external temperature and humidity sensors. Three sensors are installed in the preselector for environmental monitoring. The first is a combined temperature and humidity sensor (DTS1) which is located in a central area of the preselector assembly. The other two sensors are temperature sensors (DTS2 and DTS3) which are directly attached to the LNA and noise source, respectively, using thermally-conductive epoxy. Monitoring of these

parameters is critical to verify functionality and integrity of the sensor system and support the credibility of the RF measurements made.

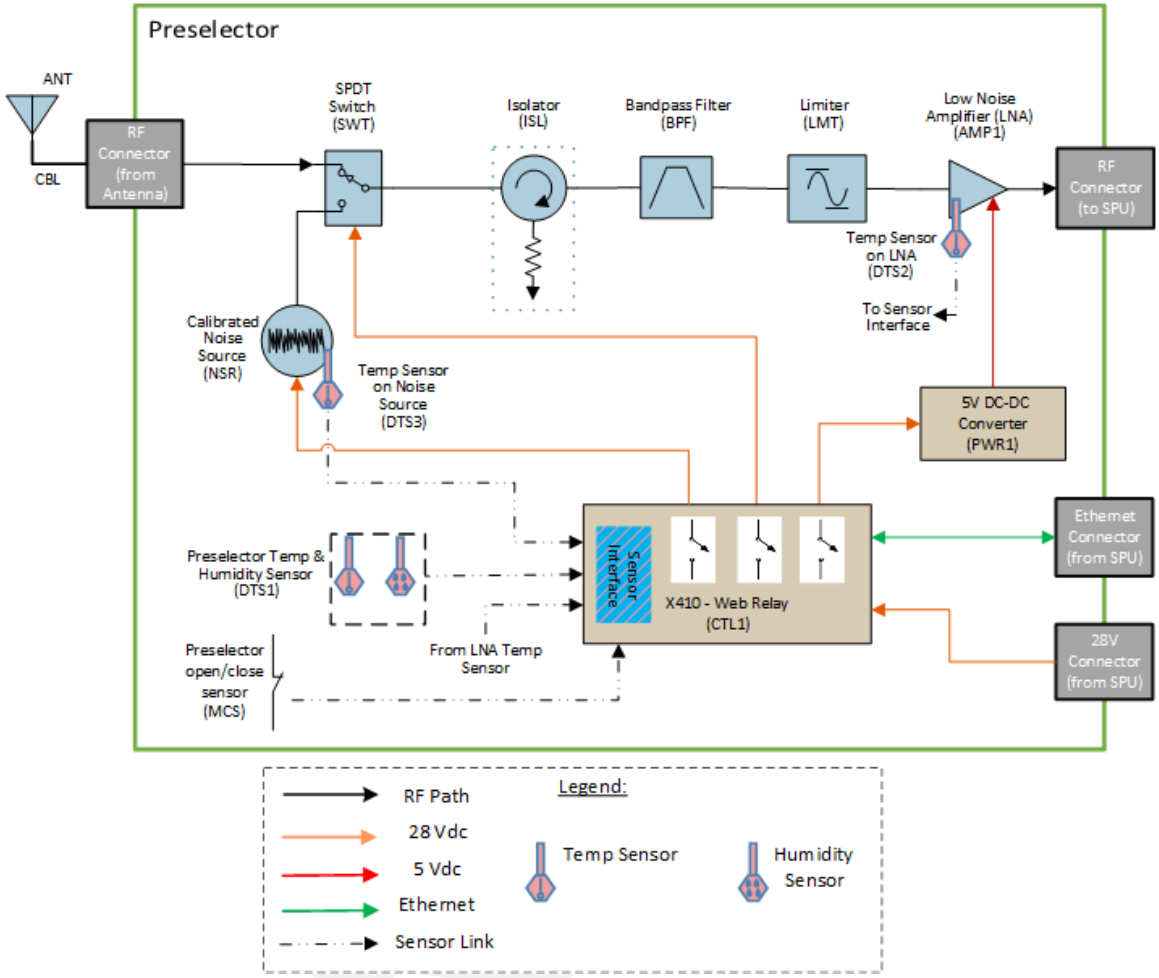

Figure 4.17: Preselector Diagram

The preselector is assembled in a 35.6 cm x 30.5 cm x 17.8 cm fiberglass-reinforced polyester weatherproof enclosure with a National Electrical Manufacturer's Association (NEMA) 4X / IP66 rating. This enclosure was selected to protect the internal electronic components from extreme environmental conditions including temperature, humidity, and solar radiation. All RF and electrical components are mounted on an aluminum base plate and interconnected with coaxial and other control and power cables. A photo of the internal assembly is shown in Figure 4.18. The bottom side of the enclosure is modified with a connector panel, shown in Figure 4.19, with feed-through connectors for RF input and output, 28 VDC power and Ethernet. An optional pole mount kit allows this enclosure to be mounted on a mast just below the antenna.

4.3.1.1 Preselector Thermal Testing

Although the preselector components do not produce a significant amount of heat, they are sealed inside a waterproof, airtight enclosure with no active cooling. As such, thermal analysis was performed to calculate the expected air temperatures inside the enclosure at the maximum power dissipation of 5.64 W and worldwide maximum ambient design temperature of 49°C. Under these conditions, the model predicted maximum air temperatures just above the components of 7.7°C without solar radiation and 10.5°C with maximum solar radiation. Figure 4.20 shows the

Table 4.2: Preselector Components

Designator	Description	Manufacturer	Part Number	Temperature Range
AMP1	Low noise amplifier, 200 MHz - 4 GHz	LiConn	LNA02004000A	-40 to +85°C
BPF	Bandpass filter, 3550-3700 MHz	K&L Microwave	13FV40-3625/ U150-O/O	-40 to +85°C
CTL1	Network-enabled controller with 4 relays and 4 digital inputs	ControlByWeb	X410-I	-40 to +80°C
DTS1	Digital temperature and humidity sensor	ControlByWeb	X-DTHS-P	-40 to +80°C
DTS2, DTS3	Digital temperature sensor	ControlByWeb	X-DTS-U	-55 to +125°C
ISL	Isolator, 2-4 GHz	CentricRF	CI2040	-40 to +70°C
LMT	Limiter, 20 MHz - 4 GHz	Fairview Microwave	FMLM2003	0 to +50°C
MCS	Magnetic contact-closure sensor	ControlByWeb	MPS	-40 to +100°C
NSR	Noise source, 10 MHz to 18 GHz	Keysight	346B	0 to +55°C
PWR1	Power supply, DC-DC converter, 28 VDC input, 5 VDC output	Acopian	28EB5E250	-20 to +71°C
SWT	Coaxial switch, single-pole double-throw DC-18 GHz	Dow-Key	401-6308	-25 to +65°C

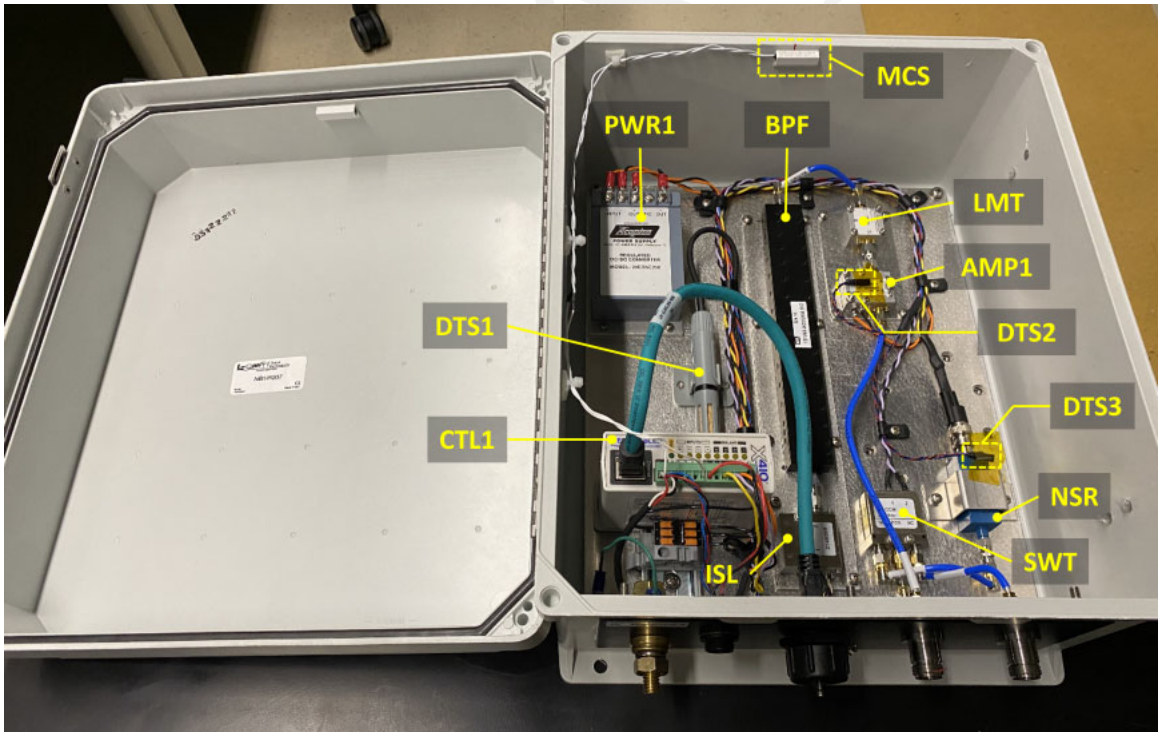

Figure 4.18: Preselector Assembly

Figure 4.19: Preselector Connector Panel

predicted air temperature just above the components with and without solar radiation. Note that the intended areas of deployment do not experience ambient temperatures that high.

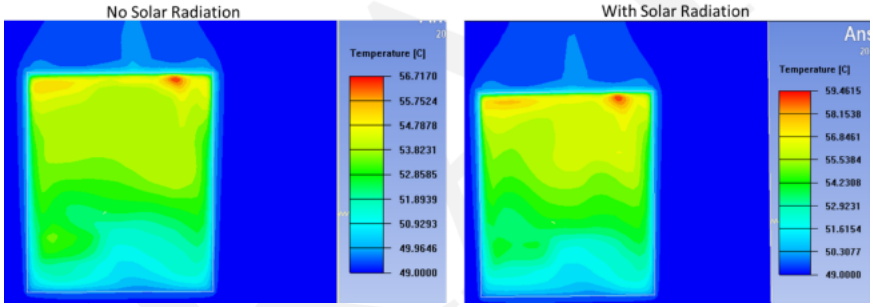

Figure 4.20: Thermal Analysis Results Without Solar Radiation and With Solar Radiation

To verify the thermal model, a preselector was installed on the roof of E Building at MITRE in Bedford, MA on 7 July 2022. Internal enclosure temperatures were monitored using three temperature sensors located inside the enclosure. The following are the locations of the sensors, as depicted in Figure 4.18

- Near center of assembly, mounted to the baseplate (DTS1)
- LNA component body (DTS2)
- Noise source component body (DTS3)

An additional sensor monitored external air temperature. However, the external sensor malfunctioned, so weather history recorded hourly at Hanscom Air Force Base in Bedford, MA during the day of the test was used instead. The enclosure was exposed to full sun for the entire test. Figure 4.21 shows the recorded temperatures.

The test started at 14:34 EDT when the enclosure door was sealed and the equipment was configured to consume its maximum power. Data was recorded for just over 3 hours. Ambient temperature peaked at 27.8°C at 15:51 EDT. The component and internal ambient temperatures peaked at about 16:32 EDT, giving a maximum rise over ambient of 8.3°C during the test.

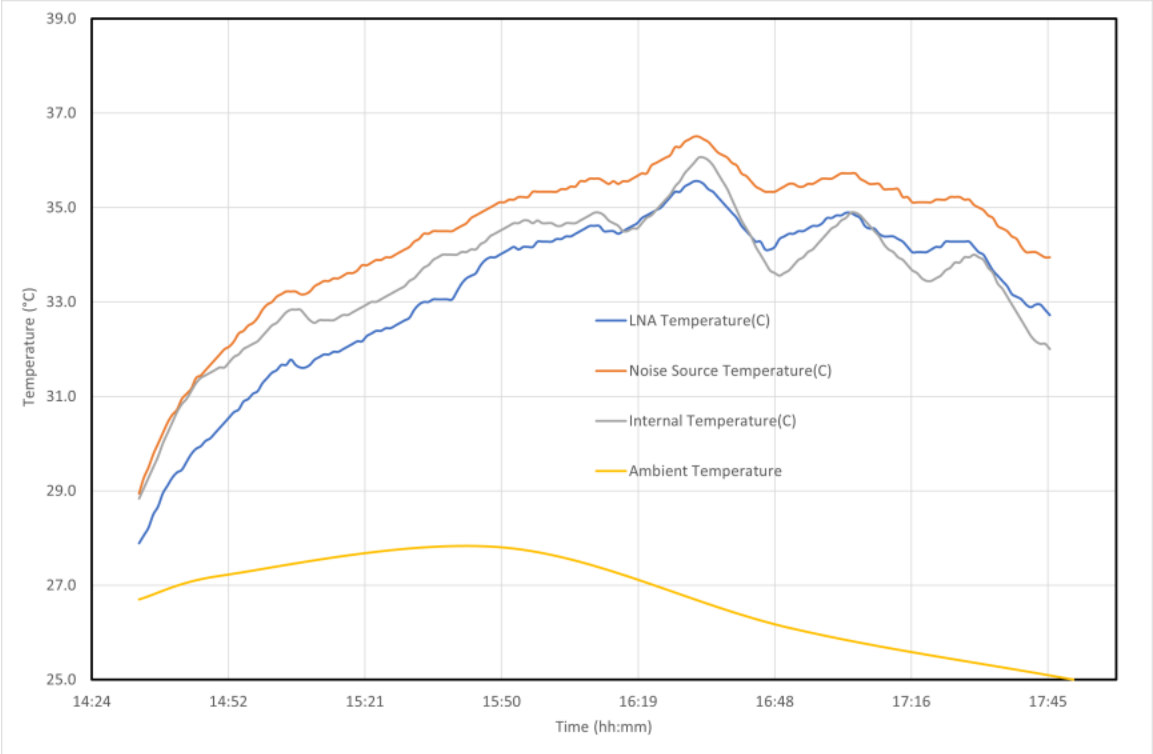

Figure 4.21: Preselector Outdoor Thermal Test Results

These results provide confidence that the equipment will operate outdoors at the maximum expected ambient temperature and exposed to solar radiation without foreseeable issues. Table 4.2 shows the operating temperature range for all components in the preselector. The component with the lowest operating temperature limit is the RF limiter (designator LMT), with a maximum operating temperature of 50°C. Considering the maximum rise over ambient temperature observed during the roof testing of 8.3°C, the preselector is expected to operate within temperature specifications up to an external temperature of 41.7°C.

4.3.2 SPU Assembly

The SPU assembly contains the remaining RF components which are not built into the preselector, namely the attenuator (ATT), second amplifier (AMP2) and RSA, as well as the embedded computer (CMP), and various system control, power management and networking components. A block diagram of the SPU is provided in Figure 4.22 and component part numbers are provided in Table 4.3.

The power management components in the RSA generate and distribute power to the entire sensor, including both internal and external components. The overall system power distribution, including these components inside the SPU, is presented in more detail in Section 4.3.3.

As with the preselector, a network-enabled relay module in the SPU (reference designator CTL2) serves a dual purpose of enabling or disabling DC power to various components and also remote monitoring of temperature and humidity. Two sensors are installed in the SPU assembly for this purpose. The first sensor is a combined temperature and humidity sensor (DTS4) located in the general vicinity of the computer and power distribution components. The second sensor is a temperature sensor (DTS5) located near the RSA.

Finally, a network switch in the SPU (reference designator NSW) creates an local area network (LAN) which enables

communication among all network-enabled components in the sensor. Specifically, these components are the network-enabled controllers in the preselector and SPU (reference designators CTL1 and CTL2, respectively), the network-enabled remote power switch (CTL3), the embedded computer (CMP), and LTE router.

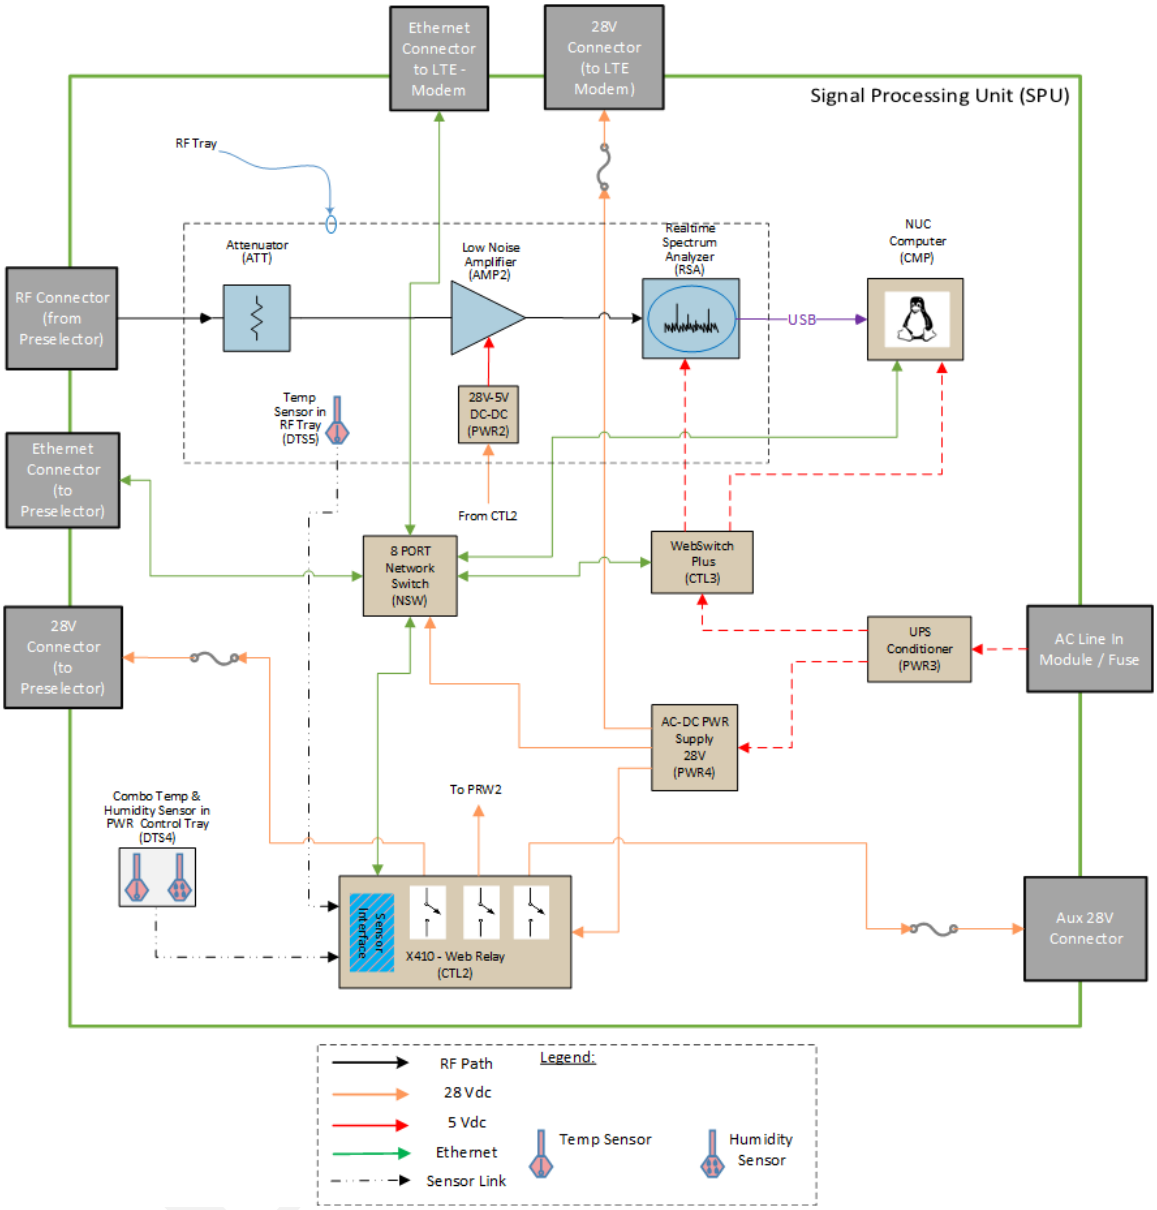

Figure 4.22: SPU Diagram

The current version of the SPU is shown in Figure 4.23. The SPU consists of several 19" rack-mount components and sub-assemblies installed into a transit case. The individual components and sub-assemblies are annotated in the figure. Starting from the top is the 1U WebSwitch module (CTL3), which is used to control AC power to the RSA and computer (CMP). Next are two custom-designed 2U trays, the RF Tray and the Power/Control Tray, separated by 1U of open space for ventilation. Next is a 2U storage drawer, and finally the 1U uninterruptible power supply (UPS) (PWR3). Note that a 7U variant of this design was also built, which is identical to the 9U shown in Figure 4.23 except that it does not have the 2U storage drawer.

Table 4.3: SPU Components

Designator	Description	Manufacturer	Part Number
AMP2	Low noise amplifier, 500 MHz - 8 GHz	Mini-Circuits	ZX60-83LN-S+
ATT	Attenuator, 6 dB, DC - 18 GHz	Mini-Circuits	BW-S6W2+
CMP	Small form factor computer, 11th Gen Intel® Core™ i7 processor, 32 GB RAM	SimplyNUC	NUC11TNHI7Z
CTL2	Network-enabled controller with 4 relays and 4 digital inputs	ControlByWeb	X410-I
CTL3	WebSwitch Plus network-enabled remote power switch and auto reboot device	ControlByWeb	XRDI-WS3P-NA
DTS4	Digital temperature and humidity sensor	ControlByWeb	X-DTHS-P
DTS5	Digital temperature sensor	ControlByWeb	X-DTS-S3C
NSW	Ethernet switch, Gigabit, 8-port	TRENDnet	TI-G80
PWR2	Power supply, DC-DC converter, 28 VDC input, 5 VDC output	Acopian	28EB5E250
PWR3	Uninterruptible power supply, 1440 VA, 1100W	Eaton	5P1500R
PWR4	Power supply, AC-DC converter, 100 - 264 VAC input, 28 VDC output	SL Power Electronics	HC28-2-A+G
RSA	Portable USB real-time spectrum analyzer, 9 kHz to 18 GHz	Textronix	RSA507A

The Power/Control Tray, shown in Figure 4.24, houses the following components. Note that the reference designators provided in the list below can be cross-referenced to Table 4.3 to obtain full part numbers.

- Ethernet switch (NSW)
- Small form factor computer with power adapter (CMP)
- 28 VDC power supply (PWR2)
- Network-enabled controller (CTL2)

On the front panel of the Power/Control Tray are connectors used for the required external interfaces. These include power and Ethernet interfaces to the preselector and LTE router, power and USB interfaces to the RF Tray, and external connections to the computer for an optional keyboard, mouse or HDMI display. The rear panel of the Power/Control Tray has two connectors for AC power inputs for the computer and the 28 VDC power supply, additional USB and Ethernet connectors, and a temperature sensor interface to the RF Tray. The Power/Control Tray front and rear panels are shown in Figure 4.25 and Figure 4.26, respectively.

The RF Tray, shown in Figure 4.27, houses the following components:

- Low noise amplifier (AMP2)
- Attenuator (ATT)
- 5V DC-DC converter (PWR2)
- Real-time spectrum analyzer with power adapter (RSA)

On the front panel of the RF tray is a type N connector that interfaces with the coaxial cable from the preselector, along with USB and DC power connectors that interface with the Power/Control Tray. On the rear panel is an AC power input connector which supplies power to the signal analyzer, as well as a temperature sensor interface connector to the

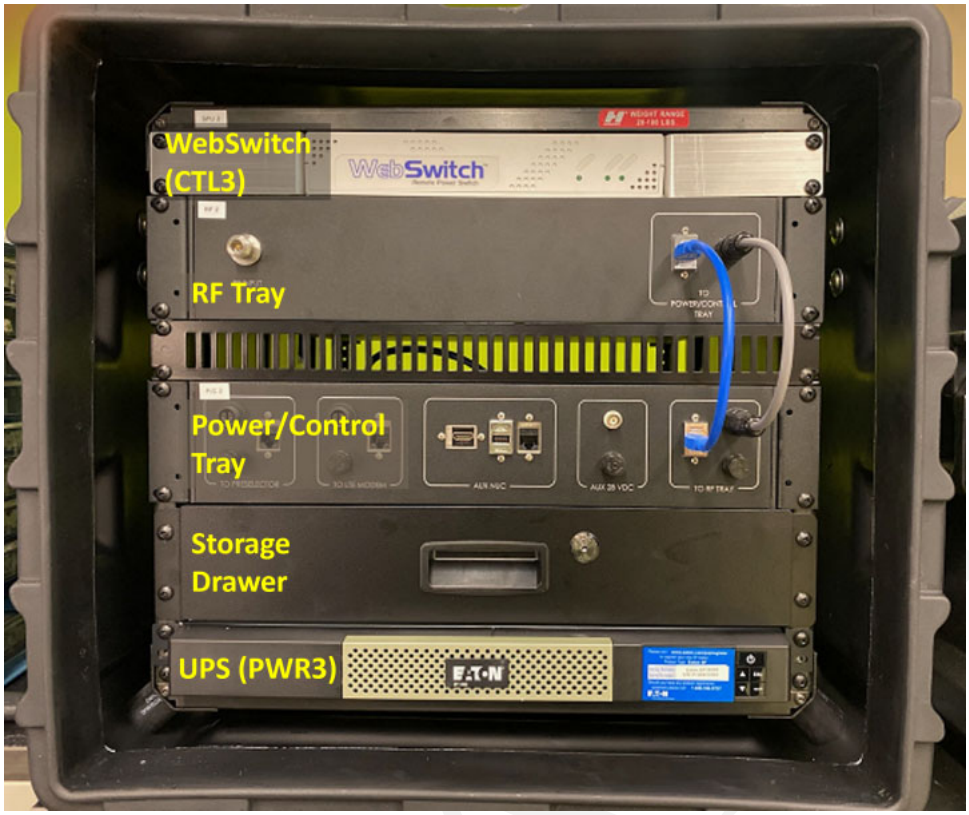

Figure 4.23: SPU Assembly (9U)

Power/Control Tray. The RF Tray front and rear panels are shown in Figure 4.28 and Figure 4.29, respectively.

Both the Power/Control Tray and the RF Tray are installed in the transit case with their tops open (i.e. no top covers installed) to allow adequate airflow and heat dissipation. The SPU must be operated in a climate-controlled area at typical ambient room temperature, and the protective front and back covers of the transit case must be removed.

4.3.3 System Power Distribution

Figure 4.30 is a diagram of the power distribution for the overall SEA sensor system. The system is powered by standard 120 VAC line power. Main power enters the system through the UPS (PWR3) located in the SPU. The UPS is capable of supplying backup power to the system for an estimated 60 minutes in the event of a power outage. From here, AC power is routed to a 28 VDC power supply (PWR4) and to a "WebSwitch" network-enabled relay module (CTL3). The WebSwitch is used to distribute AC power to the computer (CMP) and signal analyzer (RSA). Distributing power to these components through the WebSwitch enables a hard reset to be executed in the event of a system malfunction.

28 VDC power in the SPU is distributed to the LTE modem, network switch (NSW), and a second network-enabled relay module (CTL2). Controller CTL2 subsequently distributes 28 VDC to the DC-DC converter (PWR2) for the second LNA (AMP2), the preselector, and an auxiliary 28 VDC output accessed via the front panel. Relays in the X410 allow power to any of these outputs to be remotely controlled.

28 VDC power in the preselector is then distributed through another network-enabled relay module (CTL1) to the noise source (NSR) and the DC-DC converter (PWR) for the LNA (AMP1). By distributing power to these components via

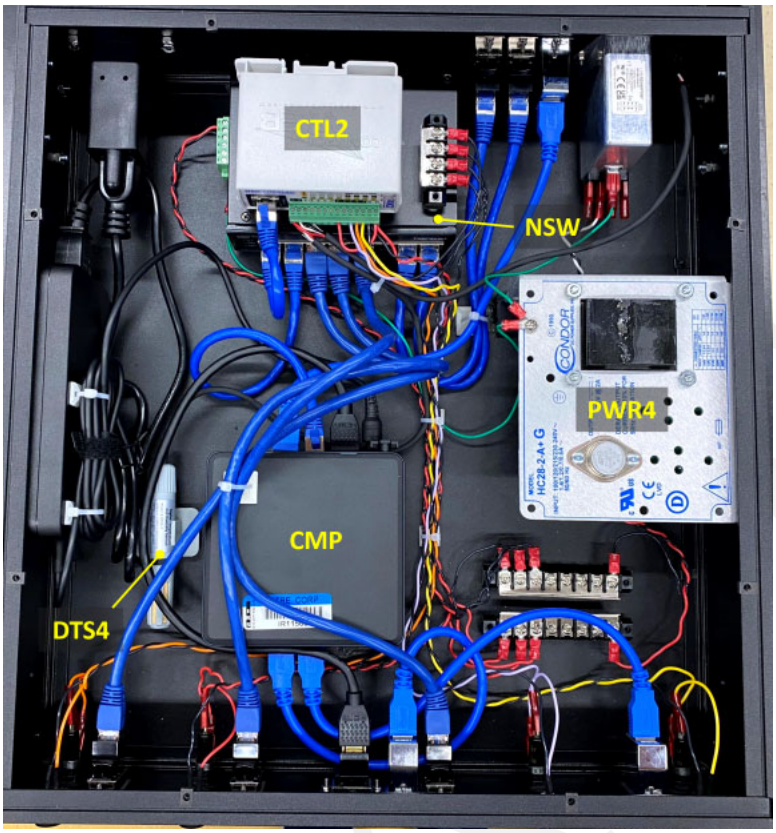

Figure 4.24: SPU Power Control Tray

Figure 4.25: SPU Power Control Tray Front Panel

Figure 4.26: SPU Power Control Tray Rear Panel

the relay module, power can be remotely enabled or disabled. The relay module also distributes 28 VDC to the RF switch (SWT). This serves as the control voltage for the switch, enabling the controller to set the switch state and select the RF input to the preselector.

4.3.4 LTE Router

Without guaranteed hardline connectivity to provide backhaul of sensor data, a wireless backhaul strategy is employed. For this, cellular routers are used as these provided the highest chance of coverage for the sensors in a variety of

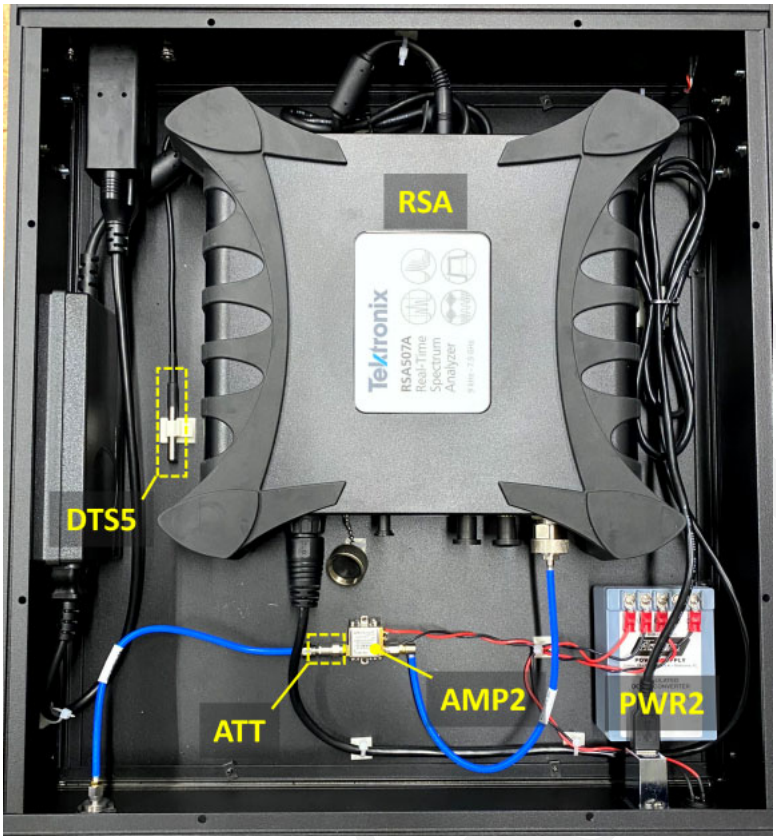

Figure 4.27: SPU RF Tray

Figure 4.28: SPU RF Tray Front Panel

Figure 4.29: SPU RF Tray Rear Panel

locations. A Digi IX20 LTE Router was chosen to enable remote access to the sensor data products as well as remote monitoring and administration of the sensor. This router provides the needed connectivity and also provides the router administration needed to provide connectivity to the individual networked sensor components (web switch, host processor, etc.). Finally, these routers are capable of being linked in a private network linked back to the management software and network.

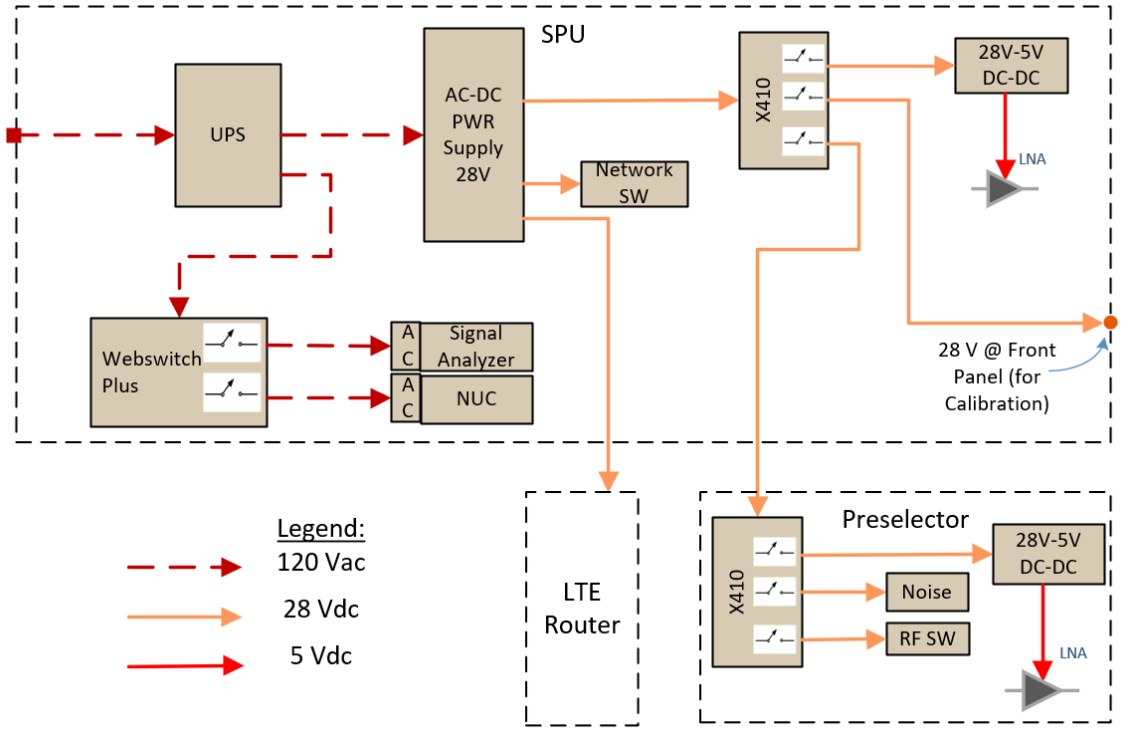

Figure 4.30: Power Distribution Diagram

4.3.5 Future SPU Enhancements

The Signal Processing Unit is currently in a final prototype version (Prototype Rev 3). National Advanced Spectrum and Communications Test Network (NASCTN) plans on design enhancements rolled into subsequent versions. The following is a summary of the current version and what is expected in NASCTN SEA sensor production versions.

- **Prototype Rev 3 Version SPU:** A ruggedized 19 inch rack enclosure (based on a 7U or 9U Pelican BlackBox) was utilized as the enclosure for this version. It is comprised of several trays containing the power, control and RF components. This was chosen for quick prototype deployment to support initial sheltered, climate controlled installations. This arrangement was chosen so that hardware and system changes can be more easily implemented as the initial design is tested and bench-marked. Digital, control, and power components are distributed across a 1U space (for the Webswitch) and in a 2U tray near the center of the stackup. RF components are located in the 2U RF tray is located just below the Webswitch space. The SPU has some limitations with its weight, size, and being essentially a laboratory grade system. To expand installation site options and include outdoor non-sheltered areas, there is a need to update the SPU design to overcome the environmental limitations, optimize its weight and size, and improve its performance. This will require an analysis of the sensors heat producing components (Computer, Spectrum Analyzer, etc) in a real world installation to estimate the amount of cooling required and an optimum component layout. Additionally, considerations on removal of humidity and SPU EMI emissions will be worked in parallel with this. The next revision (Rev 1/2) will address all these limitations.
- **Rev 1 & 2 SPU:** These future versions will incorporate various enhancements including a revised enclosure for outdoor installations featuring a sealed electronics cavity with active cooling and humidity control. This will facilitate installation in both indoors or in weather exposed outdoor environments with rain, temperature, and solar radiation. It will also feature an optimized form factor to minimize size and weight for easier installation.

Table 4.4: Sensor RF Specification Summary

	Parameter	Typical Value	Conditions	Reference
System	Frequency Range	3550-3700 MHz		Section 4.4.1
	Gain	31.5 dB	Channel 8 (3620-3630 MHz) RSA preamplifier enabled RSA reference level = -25 dBm RSA input attenuator = 0 dB	
	Input-Referred Noise Power	-98 dBm/10 MHz		Section 4.4.2
	Input 1 dB Compression Point (with CW Stimulus)	-48 dBm		Section 4.4.3
Antenna	Gain	9 dBi		Section 4.2.1.1
	Polarization	Vertical		
	Vertical 3 dB Beamwidth	11°		

4.4 Full-System RF Measurements

This section provides details on several RF measurements that were made on the fully assembled SEA sensor. Unless otherwise noted the measurements encompassed the entire RF system as shown in Figure 4.2 up to and including the RSA, but excluding the antenna. Since measurements were done on the fully assembled system, they include the impact of many small insertion losses resulting from various interconnects and front panel connectors not captured in the RF system diagram. A summary of RF performance specifications is provided in Table 4.4, and further details on these measurements are discussed in the subsections that follow. The system under test only deviated from the deployed sensors in that the 30.5 m coaxial cable between the preselector and SPU was not used. In lieu of the long coaxial cable, a short cable and attenuator with a combined insertion loss of 8.5 dB was used. This is nearly equal to the measured insertion loss of the long cable of 8.4 dB.

4.4.1 Gain and Frequency Response

The gain of the full system was measured using a signal generator to source a CW reference signal to the input of the preselector, and recording the signal level measured by the RSA. The measurement was automated so that it could be performed over a range of frequencies to characterize the system frequency response. A Rohde & Schwarz SMB100A microwave signal generator was used to sweep a signal over the frequency range 3375 MHz to 3875 MHz in 1 MHz steps. At each frequency step, the spectrum trace was acquired from the RSA and the measured signal level at the source frequency was recorded. The measured signal level and known power level of the CW input signal were then used to determine the system gain. The result is shown in Figure 4.31. This measurement shows the total system gain within the CBRS band to be approximately 30.5 dB to 31.5 dB. Outside of the CBRS band the gain rolls off closely following the response characteristics of the bandpass filter (BPF), as detailed in Table 4.1. It is important to note that this gain measurement relies on the signal level as reported by the RSA, which has a specified amplitude accuracy of $\pm$ 1.75 dB for the frequency range at which this test is performed.

4.4.2 Noise Figure

The system noise figure was measured at National Institute of Standards and Technology (NIST) following the Y-factor technique. A calibrated noise diode (NIST ID = D14B2.089, MY6140089) with an excess noise ratio (ENR) of 14.91 $\pm$ 0.05 dB was used for this measurement. The noise diode was connected to the preselector input using a 3.5 mm-SMA adapter followed by an SMA-N adapter. The RSA was configured to replicate the deployed sensors, with the preamplifier enabled, reference level set to -25 dBm and input attenuator set to 0 dB.

With a sample rate of 14 mega samples per second (Msps) (corresponding to a usable bandwidth of 10 MHz) IQ data was acquired from the RSA with the noise diode both enabled and disabled. From these IQ time series the mean power with the noise diode enabled and disabled were computed, and subsequently the Y-factor and noise figure were

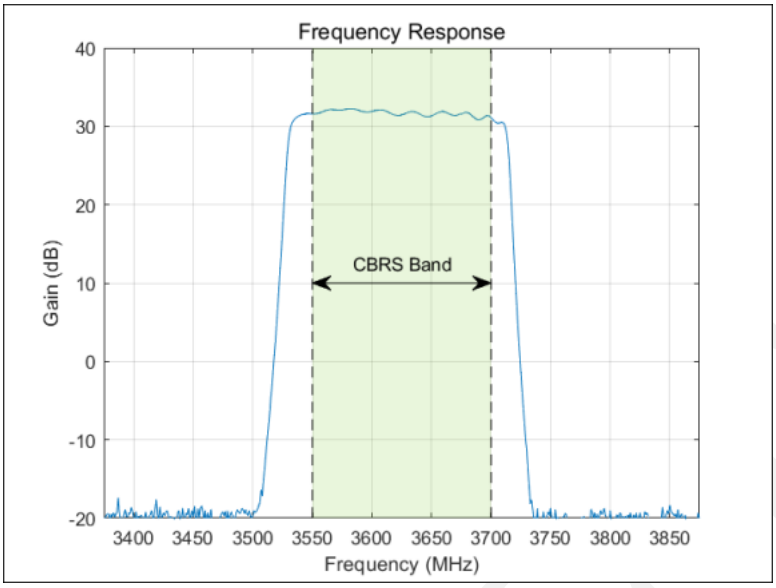

Figure 4.31: Full System Frequency Response

calculated. This measurement was repeated at 15 frequencies across the CBRS band, the result is shown in Figure 4.32.

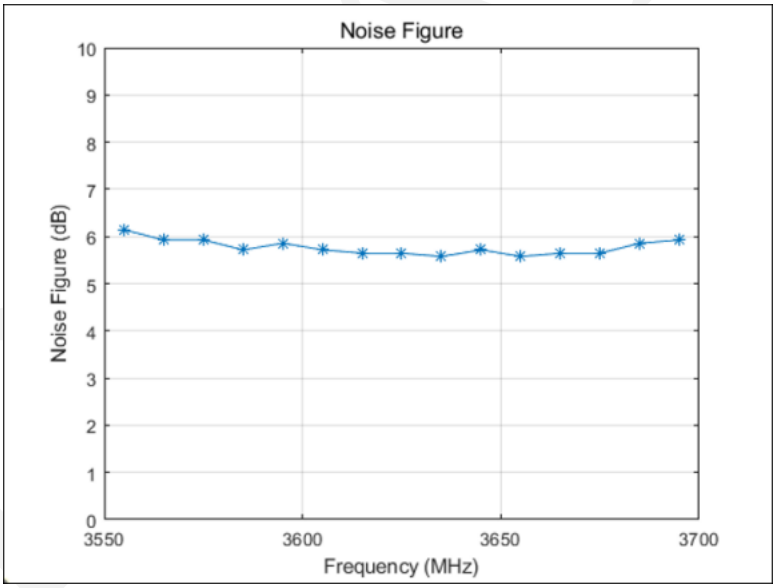

Figure 4.32: Full System Noise Figure

4.4.3 Gain compression

The input 1 dB compression point of the system was measured using a Rohde & Schwarz SMB100A microwave signal generator to source a CW signal at increasing power levels to the input of the preselector. At each input power step the signal level reported by the RSA was used to determine system gain. As the input power was increased, the point where the measured gain deviated from a constant gain by 1 dB was recorded.

This measurement was performed with an input frequency of 3625 MHz. The RSA preamplifier was enabled, a range of reference level settings were used, and the input attenuator was either set to "auto" or manually set to 0 dB. The result is shown in Figure 4.33. The deployed sensors are currently configured for a fixed reference level of -25 dBm and input attenuator of 0 dB. Under these conditions, the input 1 dB compression point was determined to be -47.7 dBm.

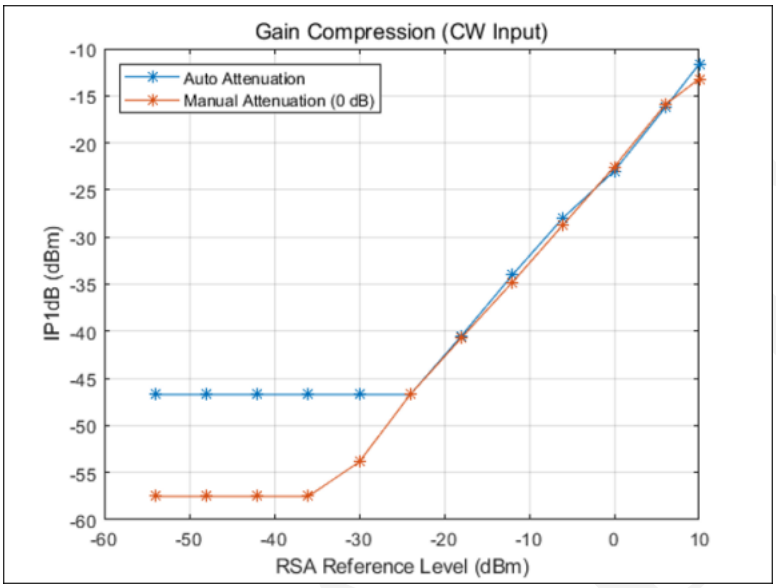

Figure 4.33: Full System 1 dB Compression Point

For most reference level settings measured, the gain remained nearly constant up until the point that the RSA reported an "ADC overrange" condition. Figure 4.9, although a measurement on the RSA only and not the full system, is representative of this type of gain compression characteristic. This indicates that the RSA, and not the RF amplifiers in system, is limiting the system 1 dB compression point. Only at the highest reference level measured, +10 dBm, and with the input power level greater than -15 dBm was compression of the RF amplifiers observed.

4.4.4 Sensor Overload

In addition to overload/compression measurements performed on the signal analyzer itself, an overload analysis was performed on the sensor system to observe the effects on adjacent channel measurements. This was performed with two input signals: a CW tone and a Gaussian noise-type signal. First, a CW tone was placed at the antenna port on the preselector and increased in power until overload was detected. Three power levels were then tested to see the data product response to the stimulus.

When the CW tone is well under the overload level as in Figure 4.34, only the channel containing the CW tone is visible above the sensor noise floor. However, once the signal increases beyond the sensor overload, the power leaks into adjacent channels. In Figure 4.35, the CW was increased beyond the overload level and the power leaked into two adjacent channels above and below the main channel. This is consistent with the native sample rate of the RSA of 56 Msps. I.E. 28 MHz above and below the center frequency corresponds to two 10 MHz channels above and below that are corrupted by an overload in the main channel.

In addition to a CW tone, a 10 MHz wide additive white Gaussian noise (AWGN) signal was used to mimic a wideband signal as a precursory stand-in for an allocation in the CBRS band. Similarly, when under the overload level, no channel bleed was observed and when the overload level was exceeded, channel bleeding was observed in the two channels above and below the channel containing the noise-like signal. The powers at which overload occurred corresponded to when the peak power of the noise-like signal exceeded the same power as a CW which caused

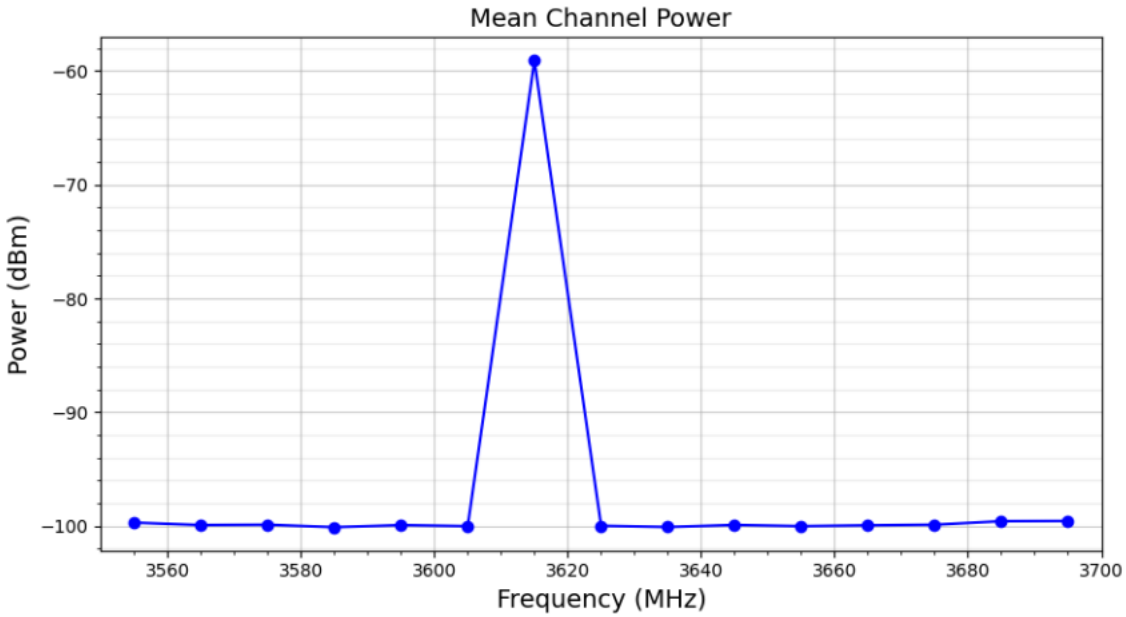

Figure 4.34: Channel power measurements with a CW stimulus well below the overload point.

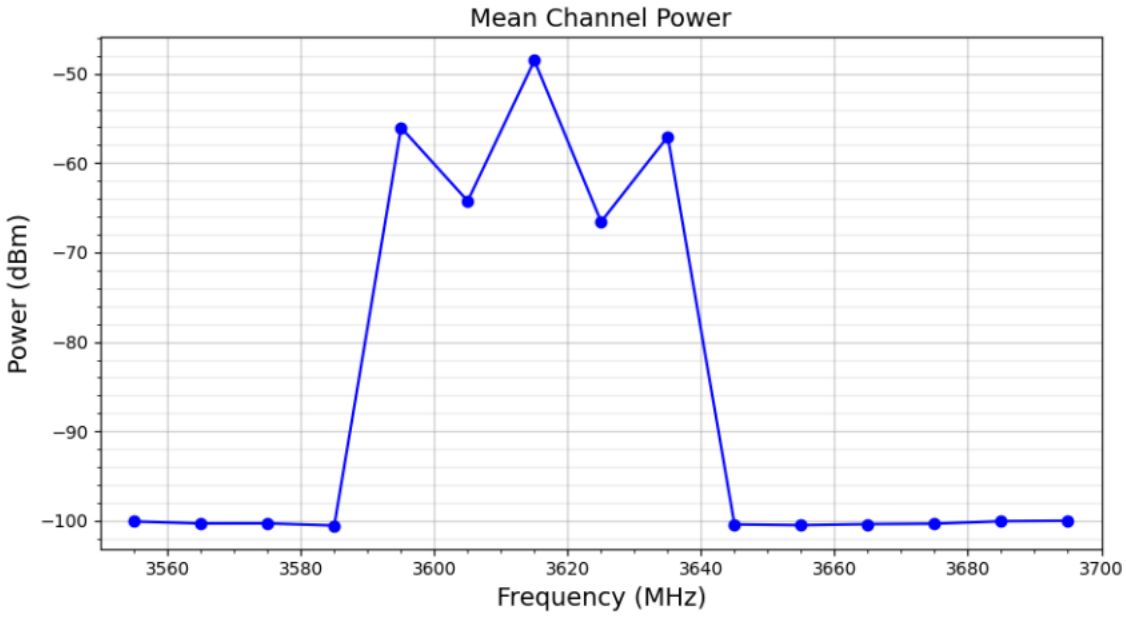

Figure 4.35: Channel power measurements with a CW stimulus above the overload level.

overload. However, unlike the CW case, when the noise-like signal power was increased further, channel bleeding into channels beyond the two channels above and below the main channel was observed as seen in Figure 4.36.

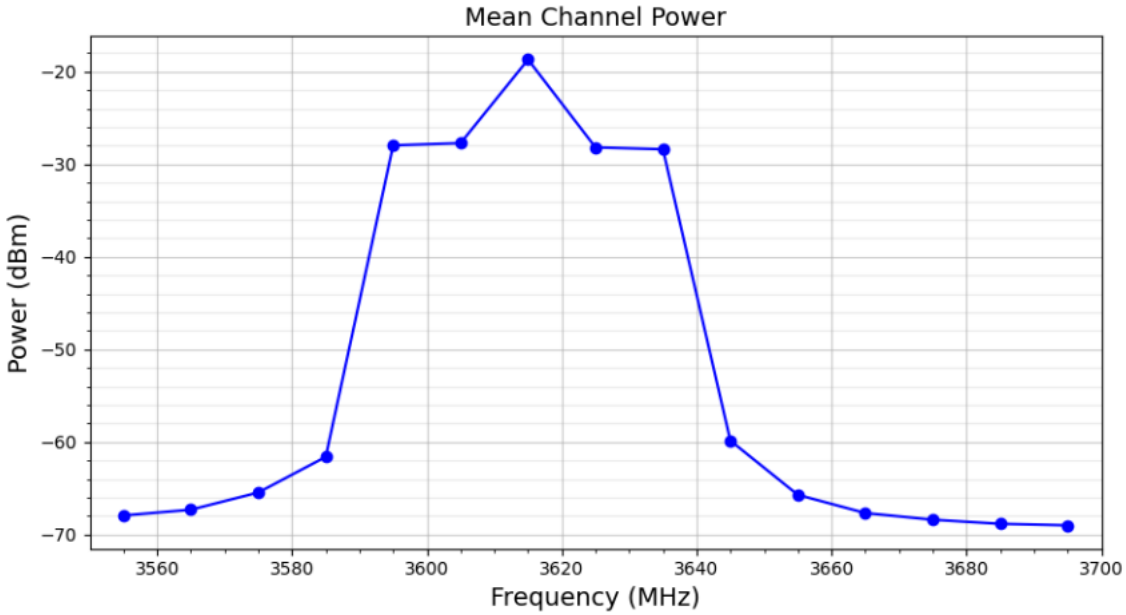

Figure 4.36: Channel power measurements with a noise-like stimulus greatly above the overload level.

Chapter 5

Sensor Management

5.1 Software

A distributed software system based on Institute of Electrical and Electronics Engineers (IEEE) 802.15.22.3-2020 [13] provides command and control over the network of sensors, and storage and retrieval for the data provided by the the sensors. As shown in Figure 5.1, the system is comprised of three main entities: an Ansible® control node, the sensor software, and the manager. Together, these entities provide automated provisioning, sensor control, and data access. The Ansible® control node automates software provisioning. Ansible® playbooks install signal analyzer dependencies,

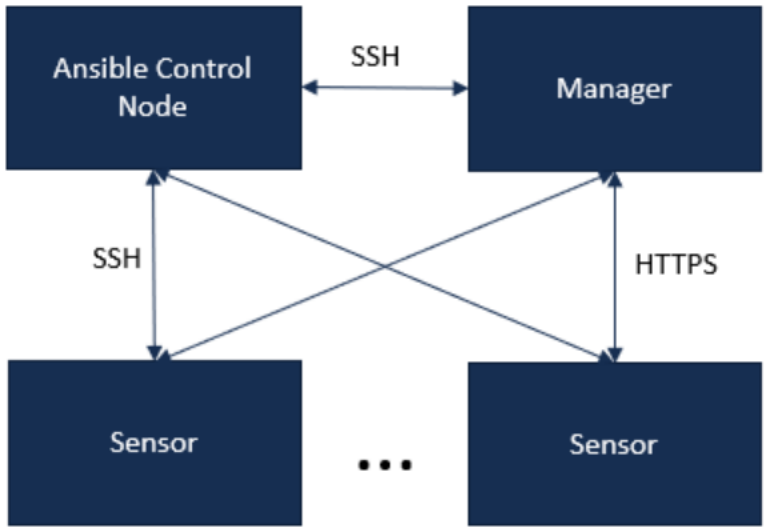

Figure 5.1: A block diagram indicating the high-level software architecture.

configuration files, sensor software, and apply the Department of Defense (DoD) Security Technical Implementation Guides (STIGs) [14] to secure the system. The sensor software, detailed in 5.1.1, establishes a common application programming interface (API) to interact with a sensor and the manager. Subsection 5.1.3 provides details on the manager software that provides command and control over the network of distributed sensor and storage and retrieval for sensor data.

5.1.1 Sensor

The sensor software was designed to provide easy-to-use remote sensor control and data retrieval regardless of the underlying hardware or sensing capabilities of the physical sensor. This is accomplished by abstracting sensing capabilities into named actions the sensor can perform. Actions may be parameterized so a single implementation may be instantiated into multiple actions. For example, Figure 5.2 provides an example configuration file for an action to acquire 10 MHz of in-phase and quadrature (IQ) data centered at 3555 MHz. There is a single software implementation of the IQ action, but duplicating this configuration file with a different name and frequency would create a new IQ action at a new center frequency.

Sensor control is accomplished through a web-based API that includes four key endpoints:

- capabilities: Get requests provide a description of the physical components of the sensor and a list of the sensing actions the sensor can perform.
- schedule: Get requests provide a list of schedules including past, ongoing, and actions scheduled in the future. Put requests allow schedules to be modified, including deactivated. Delete requests delete schedules and Post requests add new schedule entries.

```
single_frequency_time_domain_iq:
  name: 3555_IQ
  rf_path: antenna
  frequency: 3555e6
  reference_level: -25
  attenuation: 0
  preamp_enable: True
  sample_rate: 14e6
  duration_ms: 4000
  nskip: 0
  classification: UNCLASSIFIED
```

Figure 5.2: An example configuration for the an action to acquire IQ data.

- status: Get requests provide sensor status including scheduler status, location, system time, when the sensor software started, last calibration time, disk usage, how many days the sensor has been up, signal analyzer health status, and status of the preselector and switches. Preselector and switch statuses are driven by configuration files and provide healthy Boolean, the name of the switch or preselector, and a true or false status for each configured state indicating whether each state is enabled or disabled. The Sharing Ecosystem Assessment (SEA) sensor preselector is configured to include `noise_diode_powered`, `lna_powered`, `antenna_path_enabled`, and `noise_diode_path_enabled` states. The WebRelay™located within the signal processing unit (SPU) includes `rf_tray_powered`, `preselector_powered`, and `28V_aux_powered` states. The WebSwitch™includes `sigan_powered` and `computer_powered` states.
- tasks: Get requests enable the retrieval of sensor acquisitions (signal metadata format (SigMF) archive of metadata and data). Delete requests delete the data and metadata.

As shown in Figure 5.3, the web services are implemented as Django applications hosted in a Gunicorn application server and exposed via NGINX, a high performance web server. For portability and ease of deployment, the sensor software is containerized with Docker. NGINX, the PostgreSQL database, and Gunicorn along with the sensor software are each hosted in separate containers. In addition, there is a fourth Autoheal container, described further in section 5.1.4 that monitors the health of the API container and restarts it, if necessary.

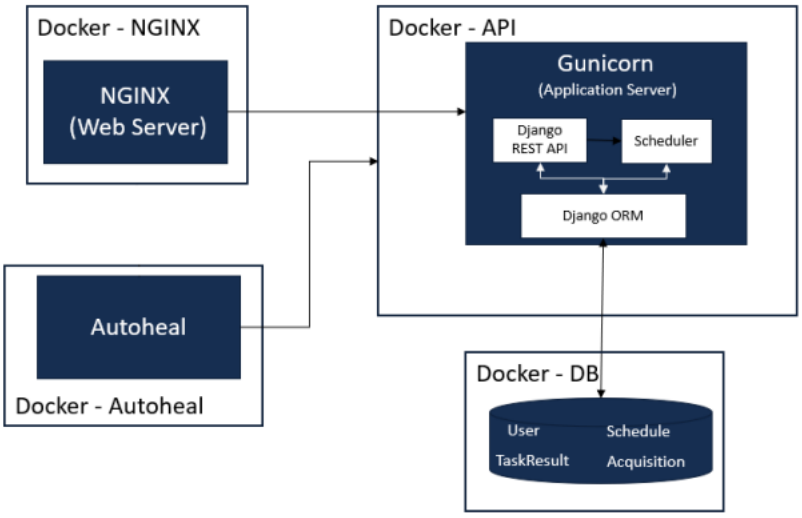

Figure 5.3: A block diagram indicating the sensor software architecture.

5.1.2 Sensor Actions

Two actions were added to enable long-term continuous monitoring within the Citizens Broadband Radio Service (CBRS) band. A calibration action, detailed in section 7.1.2.1, supports periodic calibration. In addition, a single data

product action creates the data products described in section 6.

5.1.3 Manager

The manager provides the following key capabilities:

- Authentication and Authorization: Personal Identity Verification (PIV)/Common Access Card (CAC) based authentication and authorization ensures only authorized users can access the network of sensors and any data provided by the sensors.
- Sensor Registration: Authorized users may register and remove sensors from the system. Each sensor requires a unique digital certificate that identifies the device and and all communication between the sensors and manager occur with secure two-way authentication.
- Capabilities Discovery: the sensor network may consist of heterogeneous sensors that may differ in underlying hardware and/or software capabilities. The manager allows users to interrogate sensors to discover hardware components and sensing capabilities.
- Scheduling: RF sensing capabilities are defined as distinct tasks known as actions. Actions can be scheduled on one or more sensors. A schedule includes an action and name, and may include a start time, stop time, and interval that defines the amount of time to wait between action executions.
- Data Storage and Retrieval: When a sensor executes an action it produces an acquisition, or a combination of metadata and data, and notifies the manager that new data has been acquired. The manager retrieves the data and deletes it from the sensor. All data residing within the manager may be searched for and retrieved via the Web-based graphical user interface (GUI).
- Analytics and Visualization: a well defined analytics API supports integration of custom batch and streaming analytics.

As shown in Figure 5.4, manager uses a layered architecture with well defined services that provide these capabilities.

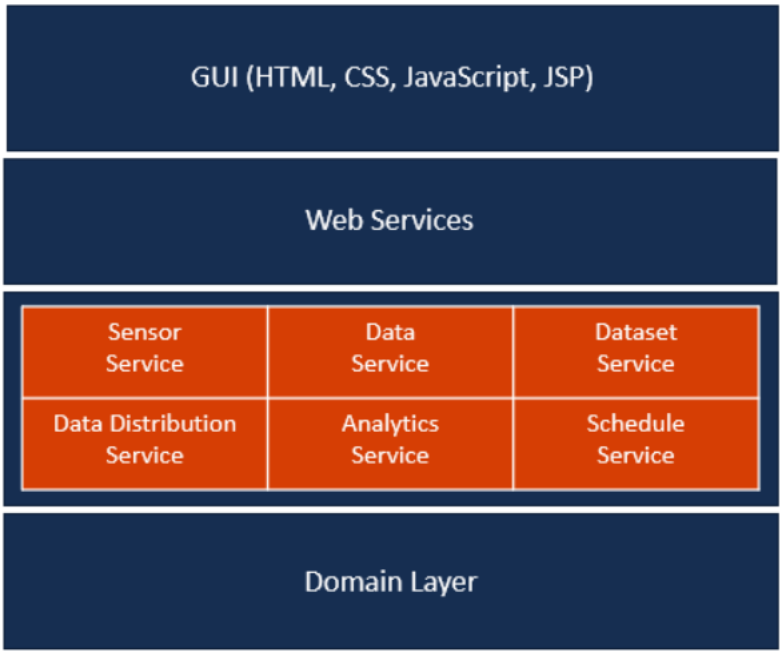

Figure 5.4: The high level software architecture diagram of the manager and its key services.

Each service implements an interface and components within the system only interact with these interfaces. In addition, service implementations are specified in configuration files and supplied to the system through dependency injection.

5.1.4 Failure Recovery

The sensor and manager software are both designed to recover from known failure conditions. The sensor software uses optional settings to indicate the WebSwitch™and states defined in the WebSwitch™to use to power cycle the signal analyzer. When the API Docker container starts, the sensor software connects to the signal analyzer. If any errors occur, the sensor software will power cycle the signal analyzer and reconnect. In addition, the sensor software employs a separate Autoheal Docker container to monitor the health of the API container. If the API container is unhealthy the Autoheal container will restart the API container. Upon restart, the API container may power cycle the signal analyzer as detailed above. At the start of each sensing action, the sensor software performs a health check. If the health check fails, the API container is marked as unhealthy.

The sensor and manager software also account for network disruptions. Every acquisition performed by the sensor is stored locally on the sensor and the sensor notifies the manager that a new acquisition is available. If the sensor is unable to notify the manager it marks the acquisition status as `notification_failed`. The manager uses a configurable background thread to check the sensors in any active schedules for any acquisitions with a `notification_failed` status and download and delete any that it discovers.

5.1.5 System Monitoring and Alerting

The data product action, detailed in 6, provides a variety of metadata such as temperature, humidity, and disk usage used to monitor the health and status of the network of sensors. The manager's analytic API allows plugins to subscribe to data provided by specific sensors and actions. A monitoring and alerting plugin subscribes to acquisitions of the data product to provide alerts if acquisitions are no longer received from a sensor or environmental conditions approach the limits of the sensor.

5.2 Network and Data Backhaul

For backhaul connectivity, each sensor is equipped with an long-term evolution (LTE) router. These routers are configured into a private LTE network with one central router bridging to the Institute for Telecommunication Sciences (ITS) network where the necessary services such as spectrum-characterization and occupancy sensing (SCOS) manager reside.

The local sensor network is managed by the local LTE router. This router handles the port forwarding necessary to branch out the same service to each device (e.g. both web relays use the same ports for their services).

This router also is part of the private LTE network with its own LTE internet protocol (IP), along with the routers from each other sensor. An additional LTE router is configured on the ITS network with routing rules and port openings to allow traffic to pass through the two networks.

Each LTE router is on a shared, data-limited plan with the limit estimated by double the estimated data accumulation (since the data is both uploaded by the sensor router and downloaded by the ITS router).

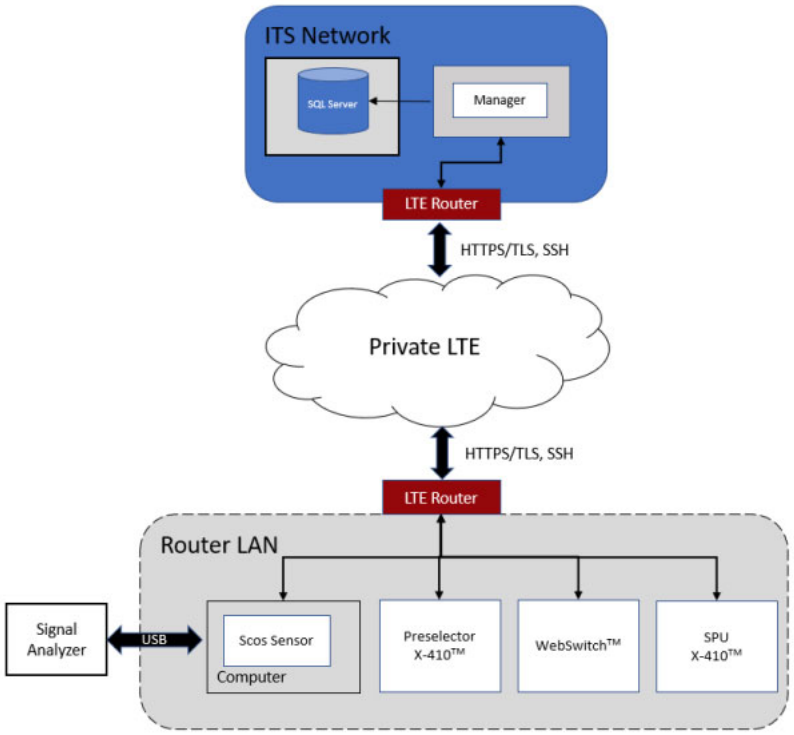

Figure 5.5: A network diagram of the connectivity between the manager and sensors.

5.3 Data Repository

The data products described in 6 are contained in SigMF [15] archives. Each SigMF archive has one or more recordings, and a recording includes a JavaScript object notation (JSON) metadata file and a dataset file containing the radio frequency (RF) measurement data.

All recordings produced by the sensors are stored in the data repository. The data repository is encapsulated within manager's data service. This provides the opportunity to change the details of the underlying repository with minimal impact on the rest of the system. The data service uses a combination of database and file system storage to store the metadata and data provided from the sensors. As shown in 5.6, a minimal subset of metadata is stored in a Structured Query Language (SQL) Server database to facilitate search and retrieval of acquisitions provided by the sensors. The full SigMF metadata file and the data file are stored on disk at the location specified by the filepath column in the database.

id	Unique id of the acquisition.
sensor	The sensor that produced the acquisition.
action	The action that produced the acquisition.
schedule	The schedule that produced the acquisition.
low_frequency	The lowest center frequency of the action.
high_frequency	The highest center frequency of the action.
latitude	The latitude of the sensor when the acquisition was produced.
longitude	The longitude of the sensor when the acquisition was produced.
start	The datetime at which the action began.
stop	The datetime at which the action stopped.
filepath	The filepath to metadata and data files.

Figure 5.6: The database table for acquisitions.

Chapter 6

Sensor Data

This chapter addresses the real-time acquisition and analysis of radio channel measurements performed by each sensor in CBRS channels.

6.1 Edge Data Payloads

Edge data payloads are the measurements produced by each SEA sensor. These data comprise a stream of power statistics and their distributions across time and frequency, together with associated metadata. These measurements serve as low-level building blocks to support analysis needed for the National Advanced Spectrum and Communications Test Network (NASCTN) SEA project. As such, a subset of the sensor data will be used in post-analysis to produce derived measurements that address the basic test questions. Edge data payloads are thus distinguished both from the lowest level data raw signal analyzer output (baseband waveforms) and from higher-level data products developed from post-analysis of aggregated payloads.

Each sensor must be automated to operate independently as a field instrument that is tailored to assess the CBRS sharing environment. This operation requires the development of automation software that is responsible for operating the signal analyzer, digital signal filtering, statistical analysis, and integration with the SCOS software framework. By orchestrating these functions, the sensors acquire waveforms from CBRS channels, apply calibration corrections to the data, and record the end results. The measurement traces from each sweep are packaged with metadata by the SCOS software for upload to centralized data storage hosted by NASCTN.

This section summarizes the design, implementation, and performance of reference software implementation of the above CBRS sensing requirements. It has been tested in the field through deployment in 3 sites. If needed, the configuration and implementation can be adapted rapidly to accommodate changing requirements in the future. As such, the reference sensing capability can be interpreted as an implementation template for the field assessments described in the test plan.

6.1.1 Reference Architecture

The sensor performs a stepped frequency sweep across the CBRS band, channelized into 10 MHz steps. In each channel, the sensor computes statistics on RF power based on IQ waveforms streamed from its signal analyzer. Within practical data and computation constraints, the goal is to generate a rich data payload which allows for flexible and in-depth analysis in aggregate.

Overall design of the edge-computed signal processing algorithm is driven by project goals subsection 1.1.3 and operating parameters of CBRS but constrained by practical implementation limitations. For example, a key goal of the project is to assess the vacate time of measured Citizens Broadband Radio Service devices (CBSDs) in the presence of a protected incumbent. There is also a maximum data throughput rate supported by the LTE backhaul, which limits the amount of data which can be sent from sensors at the edge back to the central NASCTN data repository. Finally, there are some limitations present due to the constraints of the SPU computer, such as processing and memory capacity.

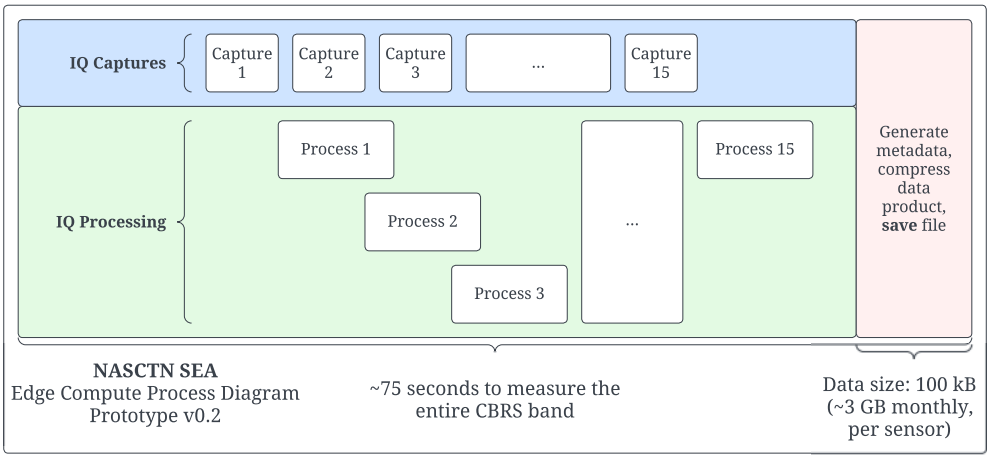

Figure 6.1: A block diagram indicating the high-level processing flow of IQ captures and data processing for a full-band edge computed data capture.

Figure 6.1 shows a high-level overview of the edge compute algorithm's processing flow and parallelization. The sensors run two primary processes in parallel: acquisition of IQ waveform captures (the blue block), and IQ sample post-processing (the green block). After the sensor records IQ samples in one channel, the post-processing of those samples begins concurrently with the next channel IQ recording. Finally, after all IQ recording and processing has been completed, capture metadata and sensor diagnostics are recorded and the results are saved (the red block). On average, the combination of these processes takes approximately 5 s in total to complete for a single 10 MHz channel using a configurable channel dwell time initially set to 4 s, with a measurement of the full CBRS band being performed in approximately 75 s. The total recording and processing is ultimately a function of the IQ recording dwell time in each channel. The dwell time itself sets a lower bound on total process time. Longer channel dwell times also result in more samples being processed into power statistics, and the computation time of these processes increases with the number of input samples.

Figure 6.2 shows the IQ post-processing which occurs for each channel after recording has completed. The IQ post-processing for each channel is implemented using a combination of serial and parallel tasks. First, the infinite impulse-response (IIR) filter is applied to the IQ waveform. Next, parallel processes compute each of power spectral density (PSD), time series power, amplitude probability distribution (APD), and periodic frame power (PFP) results. Each of these steps is described in detail below.

6.1.2 Real-Time Waveform Streaming

Working within these constraints, a signal processing algorithm was designed, to be computed at the edge on the SPU. This algorithm produces the following power statistics for the IQ waveform acquired in each channel and sweep: time series, power spectral density (PSD), periodic frame power (PFP), and amplitude probability distribution (APD). IIR filtering is used to constrain all of these power statistics to each channel. Each component of the sensor data stream, and the processing flow of the algorithm overall, are described below.

6.1.2.1 Acquisition

First, the signal analyzer is used to capture IQ samples. The analyzer is set to use a 14 MHz sampling rate, which results in 10 MHz of valid instantaneous measurement bandwidth due to the anti-aliasing filters built in to the signal analyzer itself. The center frequency is set to 3555 MHz, and a 4 s IQ capture is performed which instantaneously measures the entire 10 MHz channel from 3550 to 3560 MHz. The attenuation setting on the signal analyzer is variable dependent on the sensor location and the strength of signals observed, with the goal being to use as little attenuation as

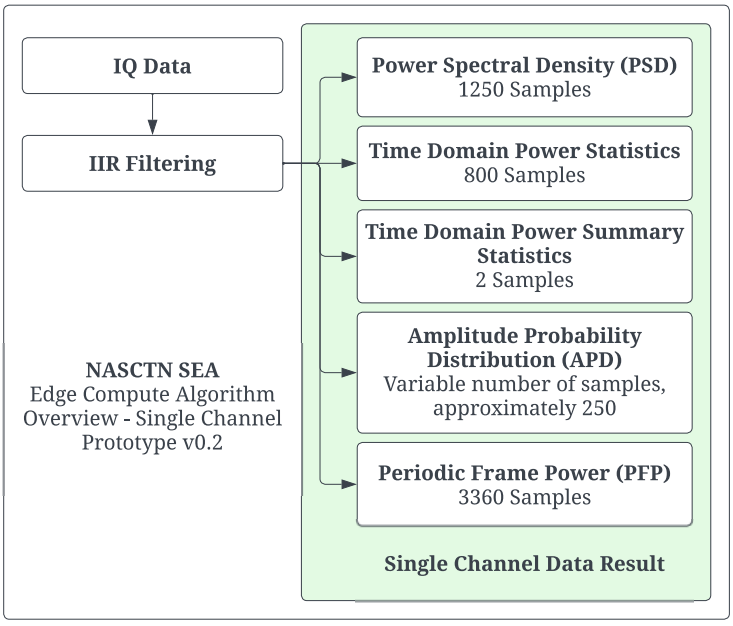

Figure 6.2: A block diagram indicating the digital signal processing flow for a single channel IQ capture within the edge compute algorithm. The IQ samples are filtered and then processed in parallel to generate the various components of the edge data payload.

Table 6.1: An overview of edge compute parameters used for prototype IQ capture and data processing.

Category	Parameter	Value
IQ Capture and signal analyzer	Record time	4 s
	Center Frequency	3555, 3565, . . . , 3695 MHz
	Preamp Enable	True
	Reference Level	-25 dBm
	Sampling Rate	14 MHz
IIR Filter	Max. Passband Ripple	0.1 dB
	Min. Stopband Attenuation	40.0 dB
	Passband Edge Frequency	5 MHz
	Stopband Edge Frequency	5.008 MHz
Time Domain Power Series	Detector Window	10 ms
PSD	discrete Fourier transform (DFT) Length	875
	Number of DFTs for Detection	64,000
	Window Type	Flat top
	Frequency Resolution Δf	16 kHz
PFP	Reference Frame Period	10 ms
	Resolution	1/56 ms ($\approx$ 17.9 µs)
APD	Power Bin Size	1 dBm

possible while avoiding signal analyzer overload on any channel. A single channel IQ capture results in a discrete-time sequence of complex-valued samples of a voltage waveform:

$$\mathrm{IQ}[n] = I[n] + jQ[n]. \tag{6.1}$$

Here, $I[n]$ is the in-phase component and $Q[n]$ is the quadrature component of the quadrature baseband signal. The sample index n ranges from 0 to 55,999,999 for a 4 s IQ capture using a 14 MHz sampling rate.

6.1.2.2 Calibration Corrections

Before any other stage of processing, the channelized IQ data is scaled using a gain factor obtained from calibrations, effectively changing the reference point of the data from the input of the signal analyzer up to the calibration point (near the antenna terminal). Y-factor calibrations using the on-board reference noise source (see Section 7.1.2.1) are automatically performed on a periodic basis in the field after sensor deployment. This routine yields a calibrated power gain value for each channel. Note that actual signal gain is frequency-dependent, but the calibration routine generates a single calibrated gain value representing the entire 10 MHz channel. If the calibrated channel power gain, in linear units, is G_{eff}, then the scaling is applied to recorded IQ samples by Equation 6.2. For details on the computation of G_{eff}, refer to section 7.1.2.1.

$$\mathrm{IQ}[n] = \frac{I[n]}{\sqrt{G_{\text{eff}}}} + j\frac{Q[n]}{\sqrt{G_{\text{eff}}}} \tag{6.2}$$

6.1.2.3 Digital Filtering

Background: In earlier analysis stages of the signal analyzer response, it was shown that the integrated hardware anti-aliasing filters provide slow roll-off in frequency response outside of the measured channel. This presented a need for additional digital filtering of IQ samples, so that power statistics and other data results can be accurately and independently reported for each 10 MHz channel. Digital filtering improves the out-of-band rejection performance for each channel power measurement when the input power level is below the sensor overload level.

Description: An IIR filter was designed which was optimized for filter stability, high passband flatness, high stopband rejection, and small transition band from passband to stopband. Due to the nature of the power statistics being reported, linear phase response was not an important consideration, so IIR filtering was chosen over finite impulse response filtering. This results in faster and more efficient digital filtering of large streams of IQ samples at the edge, when compared to an equivalent finite impulse-response (FIR) filter design. For comparison, a FIR filter designed using the Kaiser window method, with the same minimum stopband attenuation and transition width as the implemented IIR filter, would have 3909 taps. Implemented for comparison on a high-performance laptop, filtering a single 10 MHz channel IQ capture took 1.5 seconds using the IIR filter and 72.5 seconds using the comparable FIR filter.

The IIR filter was designed and implemented as a digital Cauer filter of order 12. This elliptic filter type was chosen due to its ability to minimize the transition interval between the passband and stopband for a given amount of ripple [16]. The resulting filter has a maximum passband ripple of 0.1 dB below unity gain, minimum stopband attenuation of 40.0 dB, and a transition width of 8 kHz between the passband and stopband. With such a small passband ripple, the filter approximates the response of a Type II, or inverse, Chebyshev filter. The filter is implemented as a low-pass filter, which behaves like a band-pass filter for a complex input sequence such as the stream of input IQ samples. A low-pass filter with a 5 MHz cutoff frequency has 10 MHz bandwidth filter when applied to the double-sideband complex baseband signal. The filter was analyzed in the Z-plane and found to be stable since its poles lie within the unit circle. To avoid numerical errors and potential reduced filter stability due to finite floating point precision, the

Table 6.2: The coefficients of the IIR filter used in IQ data processing

i	a_i	b_i
0	0.22001756	1
1	1.89508588	5.98460684
2	8.08369813	19.1994547
3	22.28438409	40.7912472
4	43.9358511	63.2429677
5	65.02462875	74.3311099
6	73.93117717	67.6982677
7	65.02462875	47.8732528
8	43.9358511	26.1496244
9	22.28438409	10.7528549
10	8.08369813	3.21640614
11	1.89508588	0.636398683
12	0.22001756	0.0740808688

filter is implemented using second-order sections. The filter's transfer function is defined to be:

$$H(z) = \frac{\sum_{i=0}^{12} b_i z^{-i}}{\sum_{j=0}^{12} a_j z^{-j}} \tag{6.3}$$

The filter coefficients are given in Table 6.2. The filter's amplitude and phase responses are shown up to the Nyquist frequency in Figure 6.3. Additionally, the amplitude and phase responses are shown over smaller frequency spans near the transition region in Figures 6.4 and 6.5. The transfer function is represented in the Z-plane using a pole-zero plot in Figure 6.6. Finally, the filter's step response is shown in Figure 6.7.

Uses: The filtering stage is performed after IQ samples are scaled for calibrated sensor gain, and before all other subsequent, parallelized processing stages. Therefore, the benefits of limiting the measurement bandwidth to the 10 MHz channel are carried on to each of the edge data payload components.

Calculation: The recorded IQ samples are filtered using the described IIR filter, implemented as a series of second-order filters with direct-form II transposed structure. This structure uses the minimum number of delay elements and multiplications, and implements the zeros before the poles. These qualities minimize numerical precision errors during filtering.

6.1.3 Power Statistics

6.1.3.1 Time-Series

Description: The channel power time-series quantifies the time-progression of mean and maximum channel power values across the duration of a waveform capture. In the reference implementation, peak and root mean square (RMS) power detectors are applied on short time-scales to give complementary insights about short bursts and average power, respectively. The time-resolution of the detectors, at 10 ms, is set to a single radio frame in the physical layer of cellular communication standards; insight on shorter time-scales is provided by the PFP, detailed in Section 6.1.3.3.

Background: This is a common type of measurement trace available on commercial test instruments, including signal

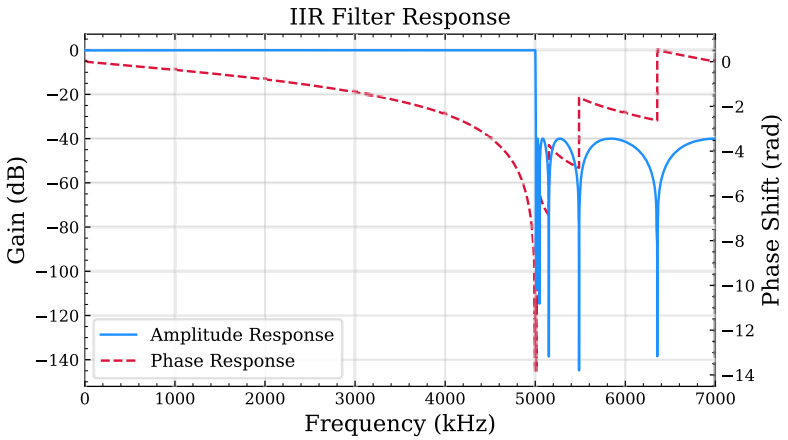

Figure 6.3: The amplitude and phase response of the IIR filter, shown up to the Nyquist frequency. The amplitude response displays a sharp cutoff at 5 MHz.

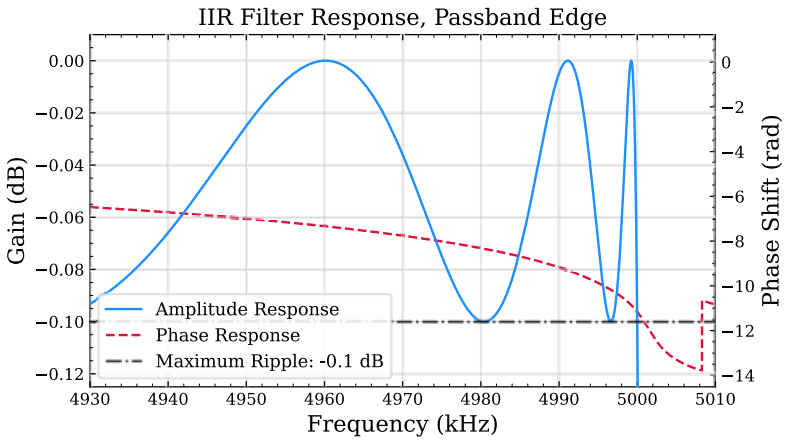

Figure 6.4: The amplitude and phase response of the IIR filter, zoomed into the region from 4930 kHz to 5010 kHz to show the passband ripple.

analyzers and oscilloscopes. These instruments sometimes refer to this measurement as *channel power*, or *power vs time*.

Uses: The time-series can be used to identify the incident time and duration for exceptionally powerful signals present in the recorded waveform. Time-series traces of power can also be used to to compute summary statistics of the time-domain characteristics of each waveform capture. The linear power average, equivalent to the RMS power detection with bin size equal to the full duration of the waveform, is the mean of the RMS power trace. The instantaneous peak power of the capture is the maximum value of the peak power trace. These are useful metrics for quantifying average and instantaneous peak power recorded in the channel for comparison with CBRS protection thresholds.

Calculation: A sequence of instantaneous power samples are calculated from IQ waveform samples by

$$P[n] = \frac{|\mathrm{IQ}[n]|^2}{2Z_0}, \tag{6.4}$$

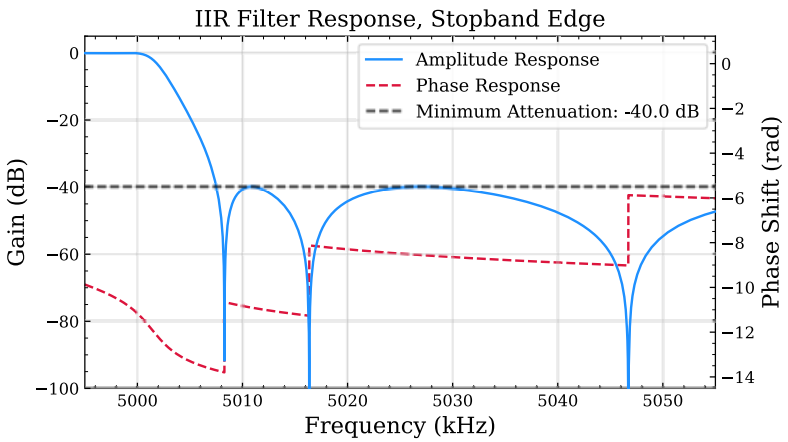

Figure 6.5: The amplitude and phase response of the IIR filter, zoomed into the region from 4995 to 5055 kHz to show the transition and part of the stopband response.

where $Z_0 = 50\,\Omega$ is the reference impedance of the sensor input. The time-series power statistics are computed for a single-channel IQ capture by the following procedure.

- Scale IQ samples by the calibrated sensor gain per Equation 6.2.
- Filter IQ samples using the IIR filter described by Equation 6.3 and Table 6.2.
- Convert IQ samples (Volts) to instantaneous power samples (Watts) using Equation 6.4.
- Reshape instantaneous power data into 400 contiguous blocks of 140,000 samples each, with each block of samples representing power samples in contiguous 10 ms blocks of time.
- *For the RMS power trace:* Take the mean of each block of samples, yielding 400 samples representing time-sequential RMS signal power on 10 ms timeframes.
- *For the peak power trace:* Take the maximum of each block of samples, yielding 400 samples representing time-sequential peak signal power on 10 ms timeframes.
- Convert both results from Watts to dBm.

Examples:

Examples of these power statistics are shown in Figures 6.8, 6.9, and 6.10. Figure 6.8 shows the results for simulated system noise. In this figure, the RMS power detector results accurately reflect the simulated system's expected noise floor. Figure 6.9 shows the results for simulated system noise and a simulated "Radar 1" signal, with its boresight pulse occurring at $t = 0.25$ s. In this figure, the simulated radar pulse is clearly visible as a spike in both traces; the RMS power trace shows a smaller spike, since the radar pulse time is smaller than the detector window of 10 ms. Note that the channel power average, $P_{avg} = 99.9$ dBm deviated by only 0.1 dBm from the value for simulated system noise without radar shown in Figure 6.8, while the instantaneous peak channel power P_{max} deviates substantially. Figure 6.10 shows the results for real-world data, in the 3590-3600 MHz channel, recorded by the prototype sensor deployed at Green Mountain Mesa in Boulder, Colorado.

6.1.3.2 Power Spectral Density

Description: The power spectral density quantifies the distribution of signal power with respect to frequency within a waveform capture. A set of power spectra are computed from the IQ waveform, each with frequency resolution $\Delta f = 16$ kHz. Peak and RMS statistical detectors are then applied to give complementary insights about short bursts and average power, respectively. The frequency resolution of the computed spectra is chosen to provide insight into

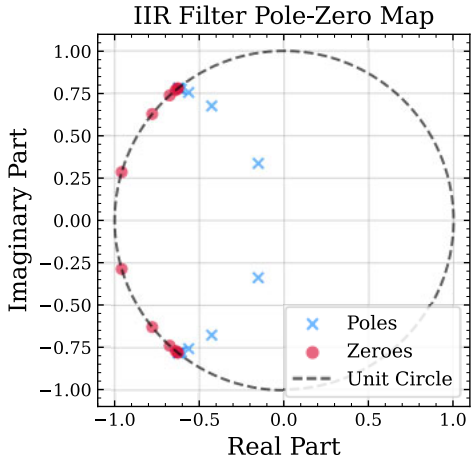

Figure 6.6: A pole-zero map, displaying the poles and zeroes of the IIR filter in the complex Z-plane. All of the poles lie within the unit circle, indicating filter stability.

the frequency distribution of the power of cellular communication signals within their resource blocks, and to allow for seamless aggregation of adjacent channel results. Windowing is used to reduce the undesirable effects of spectral leakage and scalloping loss.

Background: Representing RF signals as functions of frequency rather than time is a common practice in radio science and engineering. Viewed this way, it is possible to easily detect important signal features and characteristics, such as dominant signal frequencies or the presence of frequency-dependent signal structures. The discrete Fourier transform (DFT) provides a method for sampling the discrete-time Fourier transform (DTFT), which is equivalent within the Nyquist bandwidth to the Fourier transform. The DFT calculates the sequence $X[m]$ from an input IQ$[n]$ by:

$$X[m] = \sum_{n=0}^{N-1} \text{IQ}[n] \cdot \exp\left(-\frac{j2\pi nm}{N}\right) \tag{6.5}$$

The corresponding sequence of frequency values $F[m]$ is determined using the sampling rate F_s and center frequency f_c pertaining to the input sequence IQ$[n]$ by:

$$F[m] = \frac{mF_s}{N} + f_0, \quad 0 \leq m \leq N-1 \tag{6.6}$$

The DFT is commonly implemented in practice using the fast Fourier transform (FFT) algorithm, which allows for the DFT to be computed much more efficiently than by using Equation 6.5 directly ($\mathcal{O}(N \log N)$ vs. $\mathcal{O}(N^2)$) [17].

One characteristic of the DFT is spectral leakage, which occurs due to the implicit assumption in the DFT that the input signal is exactly periodic over the length of the input. For any frequency component of the input which is not an exact multiple of Δf, this assumption fails. Since all recorded signals cannot be expected to be constrained in this way for any choice of Δf, leakage is practically unavoidable in DFT-based processing for this measurement application. Windowing is therefore used to reduce the effects of leakage. Windows are periodic functions which go to zero at

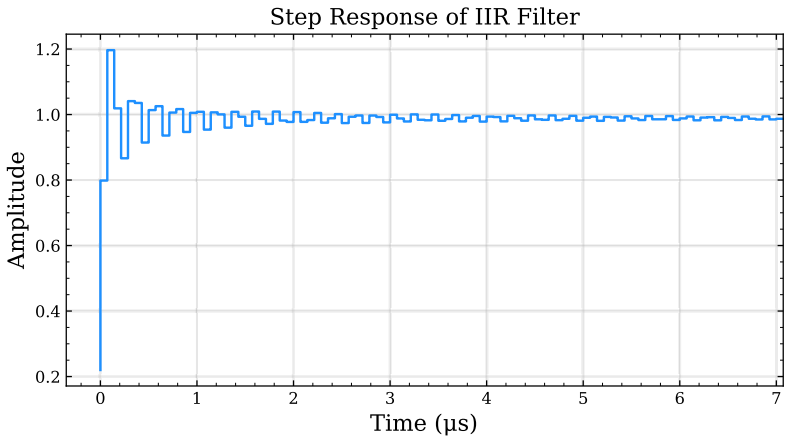

Figure 6.7: Plot showing step response of the IIR filter, demonstrating the initial overshoot and subsequent ringing effects.

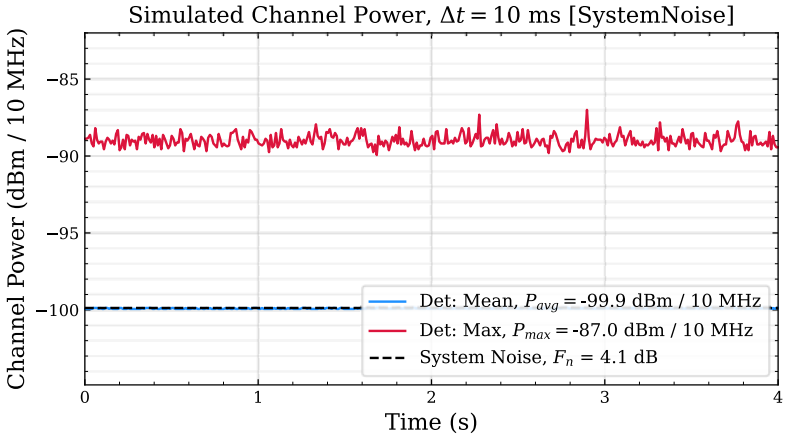

Figure 6.8: Mean and max. detected power statistics over a 4 s interval for simulated system noise in a 10 MHz channel, with simulated channel noise figure $F_n = 4.1$ dB. The predicted noise floor is shown as a black dashed line.

the first and last sample of the DFT input, satisfying the periodicity assumption of the DFT. Many types of windows exist, each with various advantages and disadvantages. The flat top window was chosen due to its good performance in reducing spectral leakage as well as its optimal performance in reducing scalloping error. The flat top window has worst-case scalloping error of 0.01 dB, while demonstrating maximal flatness in its main lobe [18]. The flat top window sequence is a 5th-order cosine function, defined by Equation 6.7 and Table 6.3.

$$w[n] = \sum_{l=0}^{L-1} A_l \cdot \cos\left(\frac{2\pi l n}{2N+1}\right) \tag{6.7}$$

The application of the window inherently reduces the amplitude and energy of the signal to which it is applied. In order to accurately measure the power spectral density at the calibration terminal, the window energy correction factor is calculated from the window by Equation 6.8.

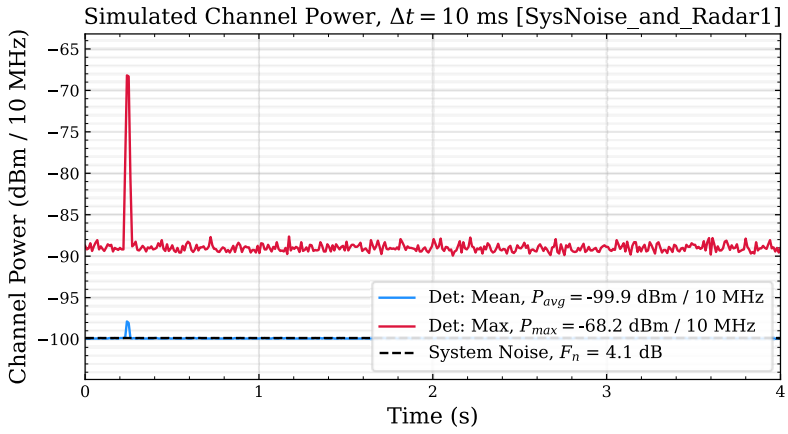

Figure 6.9: Mean and max. detected power statistics over a 4 s interval for simulated "Radar 1" and system noise in a 10 MHz channel, with simulated channel noise figure $F_n = 4.1$ dB. The predicted noise floor is shown as a black dashed line.

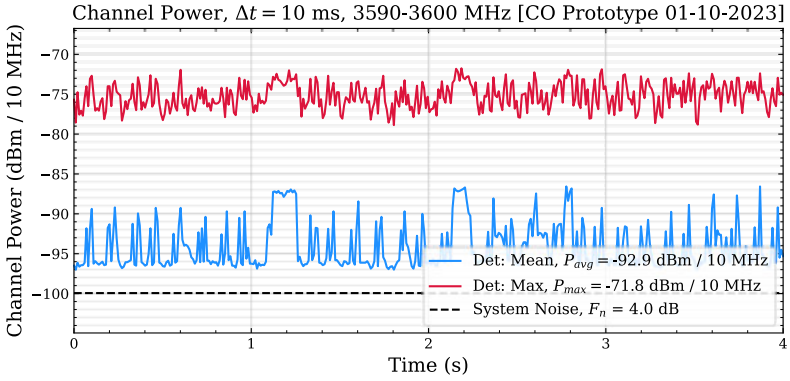

Figure 6.10: Mean and max. detected power statistics in the 3590-3600 MHz channel over a 4 s interval of a real-world signal recorded on January 10, 2023 by the prototype sensor at Green Mountain Mesa, Colorado. The system's calibrated noise floor in the channel is shown as a black dashed line.

$$\mathrm{ECF} = \sqrt{\frac{N}{\sum_{n=0}^{N-1} w[n]^2}} \tag{6.8}$$

Calculation: The power spectral density statistics are computed for a single-channel IQ capture by the following procedure. Note that the net effect of the various scaling factors applied is to remove the effects of windowing, normalize the FFT, and account for the equivalent noise bandwidth (ENBW) of the DFT bins. The results are the mean and maximum power spectral densities, with sample amplitudes in dBm/Hz, referenced to the calibration terminal.

- Scale IQ samples by the calibrated sensor gain per Equation 6.2.
- Filter IQ samples using the IIR filter described by Equation 6.3 and Table 6.2.
- Scale IQ samples for the amplitude loss due to windowing, by multiplying each sample by the window energy correction factor defined in Equation 6.8.

Table 6.3: Coefficients for the flat top window used in PSD statistics processing.

A_0	0.21557895
A_1	-0.41663158
A_2	0.277263158
A_3	-0.083578947
A_4	0.006947368

- Reshape IQ samples into 64,000 contiguous blocks of 875 samples each.
- Apply the window by multiplying each block element-wise by the flat top window sequence defined by Equation 6.7 and Table 6.3.
- Compute the DFT of each block (Equation 6.5 with $N = 875$) using the FFT algorithm.
- Convert the DFT result amplitudes (Volts) to power (Watts) using Equation 6.4. This step include scaling to account for RF-to-baseband downconversion.
- Scale power values to a power spectral density (W/Hz) by first dividing by the sampling rate F_s, then dividing by the DFT length N.
- Sort the PSD samples by their corresponding frequencies.
- Apply statistical detectors bin-wise; that is, the i-th sample in a detector result is determined from the i-th samples of all 64,000 PSD results.
- *For the mean trace* Take the mean of all PSD results, bin-wise.
- *For the maximum trace* Take the maximum from all PSD results, bin-wise.
- Convert both PSD results from Watts/Hz to dBm/Hz.
- Discard the first and last 125 samples from both PSD results, which represent frequency ranges outside the measured channel.

Examples:

Examples of these power statistics are shown in Figures 6.11, 6.12, and 6.13. The values P_{avg} are computed by summation of the mean PSD results, and are shown to agree with those values calculated from the time series results shown in Figures 6.8, 6.9, and 6.10 respectively, which were generated from the same data. Figure 6.11 shows the results for simulated system noise. In this figure, the mean PSD detector results accurately reflect the simulated system's expected noise floor. Figure 6.12 shows the results for simulated system noise and a simulated "Radar 1" signal centered on the simulated baseband. In this figure, the main lobe and some sidelobes of the simulated radar signal appear as wide peaks in the maximum PSD result. The mean PSD does not indicate the presence of radar due to the very brief radar pulse time. Figure 6.13 shows the results for real-world data, for the entire CBRS band, recorded by the prototype sensor deployed at Green Mountain Mesa in Boulder, Colorado. This figure shows an example of multi-channel data aggregation: the data shown was computed from 15 separate, channelized IQ recordings. The choice of $N = 875$ allows for this aggregation with constant sample spacing, even across channel boundaries.

6.1.3.3 Periodic Frame Power

Description: We developed the periodic frame power measurement for this project to help differentiate between transmissions from CBSDs and others, including pulsed radar, for forensic application Spectrum Access System (SAS) activation events. The underlying approach is to capture periodic channel occupancy characteristics that are fixed parameters of these different types of users.

Background: Past monitoring campaigns support the idea that channel occupancy and power can be used to help classify between downlink and uplink. Downlink transmissions tend to exhibit high channel occupancy, and typically broadcast in almost every frame [19]. In contrast, uplink signaling, which often consumes battery power in mobile devices, is optimized to *minimize* transmission time; as a result, measurements of uplinks showed dramatically

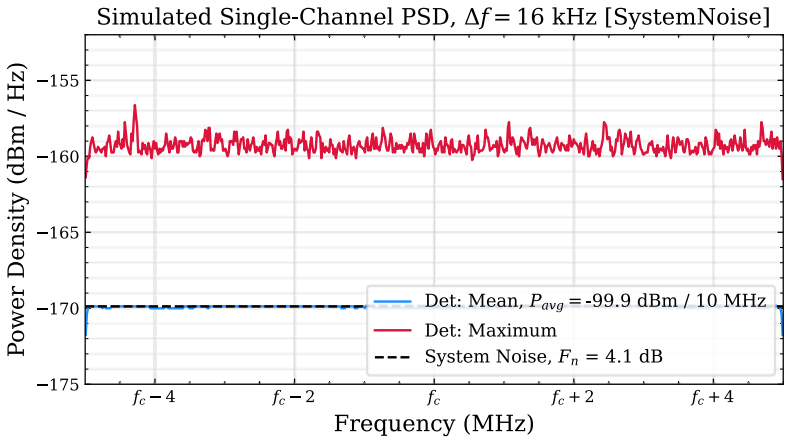

Figure 6.11: Mean and max. detected power spectral density for simulated system noise in a 10 MHz channel, with simulated channel noise figure $F_n = 4.1$ dB. The predicted noise floor is shown as a black dashed line.

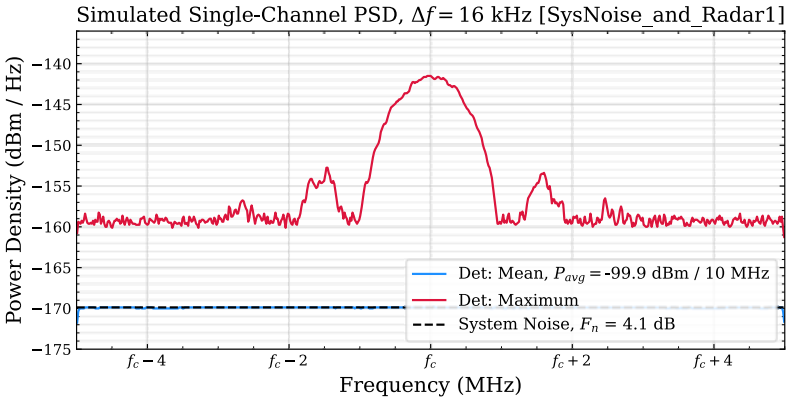

Figure 6.12: Mean and max. detected power spectral density for simulated "Radar 1" and system noise in a 10 MHz channel, with simulated channel noise figure $F_n = 4.1$ dB. The predicted noise floor is shown as a black dashed line.

lower occupancy rates, sometimes by factors as large as 10^6 [19]. These factors, together, lead to the hypothesis that aggregate power received by a SEA sensor strongest and least variable during downlink transmissions.

Each of the most widely-discussed types of incumbent and entrant users in CBRS have fixed and clearly defined periodic occupancy characteristics.

- **5G cellular** CBSDs operate time division duplex (TDD) networks with frame duration 10 ms. A base station subdivides the frame into a schedule of interleaved downlink (base station) and uplink (user equipment) time allocations. The schedule of allocations has time resolution equal to 1 symbol, which in the CBRS band may have duration 1/56 ms, 1/28 ms, or 1/14 ms [20]. Cellular standards define a large number of different such interleaved schedules that may be selected by the base station. These may remain fixed during normal operation of the base station.
- **4G cellular** CBSDs operate TDD frames with duration fixed at 10 ms by 3rd-generation partnership project (3GPP) standards [21]. A base station subdivides the schedule into a schedule of interleaved downlink and uplink time allocations. These allocations have time resolution of 1 ms. Only 7 different duplex scheduling patterns are supported, and remain fixed during normal operation of a base station. Of these,

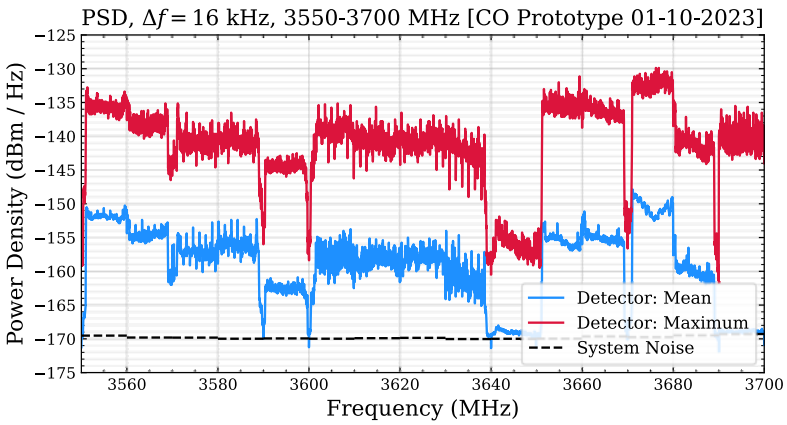

Figure 6.13: Mean and max. detected power spectral density for real-world signals in the CBRS band, with calibration-predicted channel noise floors shown as black dashed lines. The data was recorded on January 10, 2023 by the prototype sensor at Green Mountain Mesa in Colorado.

- **IEEE 802.16 (WiMaX)** CBSDs operate TDD networks. A network operates in a fixed frame duration, which must be one of {2, 2.5, 4, 5, 8, 10, 12.5, 20} ms [22]. Of these, equipment is only required to support 5 ms for WiMaX Forum certification, making this a common assumption for the actual frame duration in these networks. Recent standards also define a *superframe* as a sequence of 5 ms frames, reinforcing 5 ms as the base periodicity constant.
- "**Radar 1**" refers to one of the simulated pulsed radar systems that is used to test CBRS Environmental Sensing Capabilitys (ESCs) according to [23, Table 1]. The signal in these tests has pulse repetition interval close to 1 ms, and pulse width between 0.5 µs and 2.5 µs.

In each of these cases, a fixed division period is defined that is further subdivided to determine transmission timing. When transmissions repeat occupancy patterns within these periods, we have an opportunity to simplify the short-time features of the channel power by characterizing their typical behavior.

The key to performing power measurements of these characteristics is to define a *reference frame* that is a common multiple of all division periods under study. For our reference implementation, we selected 10 ms (Table 6.1), since it is the least common multiple of 1 ms, 5 ms, and 10 ms. Pointedly, this means that useful structure will not be visible for transmissions with aperiodic occupancy characteristics, or those which are not close to integer divisible from 10 ms.

Calculation: The calculation of periodic frame power begins with the time-series of instantaneous power samples P_n calculated from equation (6.4), and proceeds as follows:

1. Partition the time-series into M contiguous *reference frames*, following the reference frame duration (Table 6.1). For a waveform capture length of 4 s, this produces $M = 400$ frames per waveform.
2. Partition each frame into N power bins. The time-duration of these bins should not exceed the shortest possible transmission occupancy event. The reference duration of the power bins corresponds with the shortest possible duration of a single symbol defined for CBRS in 5G cellular standards: including the cyclic prefix, this duration is approximately 17.9 µs. This results in $N = 560$ power bins per reference frame.
3. Apply a power detector to evaluate channel power in each power bin. The reference implementation evaluates both peak and RMS power detectors, which are reported as separate measurement traces.
4. Compute statistics across n^{th} power bin of each frame. The resulting frame power measurement trace has M samples, each of which has been computed across the N frames. The reference implementation records {min, mean, max} for each of the peak and RMS power detectors.

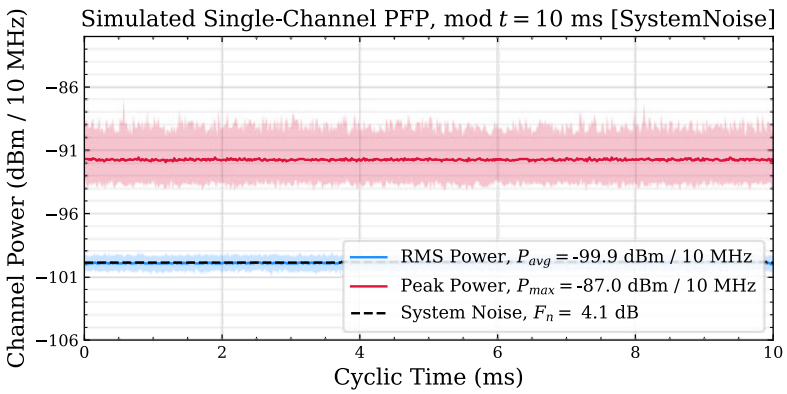

Figure 6.14: Peak and RMS-detected PFP results for simulated system noise in a 10 MHz channel, with simulated channel noise figure $F_n = 4.1$ dB. The predicted noise floor is shown as a black dashed line.

The reference frame duration should also be an integer multiple of the waveform sampling period.

Examples: The interpretation and usage of PFP is most easily conveyed by concrete example. To this end, we start with some simple simulation data, and then real field data that were observed with prototype sensors.

A simulation of the simple case of a quiet channel is illustrated by Figure 6.14, composed entirely of noise power samples. With an assumed noise figure of 4.1 dB, most of this power produced inside the sensor receiver. The solid lines represent the mean of the RMS power detector (blue) and peak power detector (red) readings; the average at each point is evaluated across the $M = 400$ data points evaluated at the same "cyclic time," which is the time elapsed into each 10-ms block. The shaded envelopes around each trace bound the extrema across the same M data points. Because the expected value of the thermal noise power in the receiver is constant, there is no time structure here except for random noise variability. Notice that the extrema for the RMS power detector are particularly close together, because the averaging detector has reduced the standard error of the noise power measurement.

Adding simulated "Radar 1" activity to the thermal noise produces the PFP illustrated by Figure 6.15. The pulse repetition interval of the waveform, which is very close to 1 ms, repeats almost exactly 10 times per frame. As a result, peaks occur close to the same cyclic time elapsed in each frame, at least relative detector resolution of 17.9 ms. The periodic characteristic is then very clear by inspection of the plot. Meanwhile, the RMS detector still shows the noise floor at the same power level as without the pulses, unchanged from Figure 6.14.

In order to illustrate CBSD activity, we move to examples of field data, since it is the duplexing dynamics cellular and WiMaX systems are proprietary and difficult to simulate. One such example is shown by Figure 6.16, taken from the prototype sensor in Colorado. In contrast with the clear pulse train in Figure 6.15, the periodic response here is divided into three regions: a steady broadcast at the highest power level, a more variable period that includes a pedestal about 15 dB lower, and a quiet period consistent with the system noise power. Further study and comparison against well-characterized CBSD traffic may in the future permit more specific conclusions to be drawn about the nature of the data networks based on these plots. However, the distinction against pulsed radar is clear.

Uses: As a culminating application of PFP is as a forensic tool to investigate CBRS DPA potential activation events. A real-world example is illustrated by Figure 6.17, for an event observed by the prototype sensor deployed at the Norfolk International Terminal (NIT) site in Virginia.

In order to start with a view across a longer time period, simple statistics were on the PFP trace in each capture

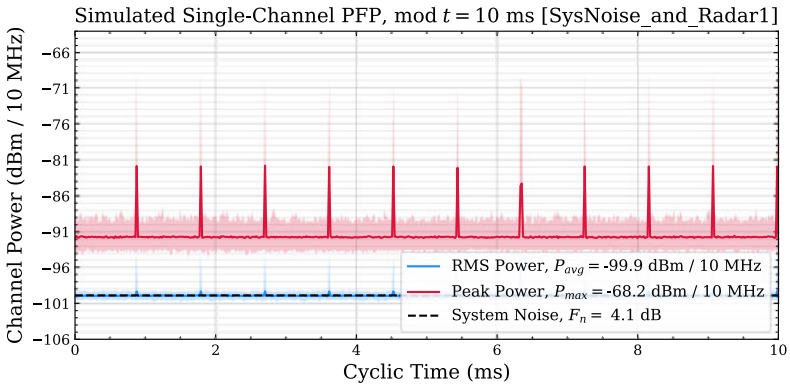

Figure 6.15: Peak- and RMS-detected PFP measurements for simulated "Radar 1" and system noise in a 10 MHz channel, with simulated channel noise figure $F_n = 4.1$ dB. The predicted noise floor is shown as a black dashed line.

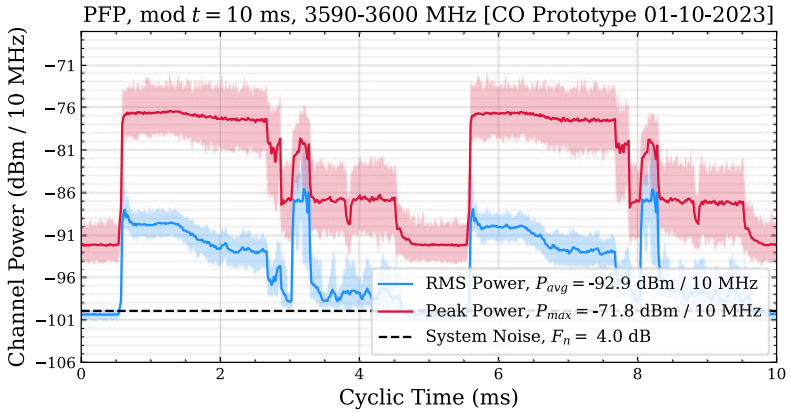

Figure 6.16: Peak and RMS-detected PFP results for real-world signals in the 3590 MHz – 3600 MHz channel, with the calibration-predicted channel noise floor shown as a black dashed line. The capture was recorded on January 10, 2023 by the prototype sensor at Green Mountain Mesa in Colorado.

for 16 minutes in Figure 6.17(a): a "capture median" (the median of average RMS detector readings), and "capture max" (maximum of the peak detector readings). Just after 7 min time elapsed, a capture maximum power reading has increased by about 10 dB, without significant change in the capture median. This suggests that a pre-existing signal has continued while short bursts have become present. Then, near 10 min elapsed, the capture median falls to the receive system noise level: for most samples in these captures, there is no detectable activity in the channel.

Figure 6.17(b) details PFP captures at the 3 key points marked along sub-figure (a). First, at 5.7 min., we see the envelope of ongoing TDD CBSD activity in the channel. By 8.6 min., a pulse train has emerged that is superimposed on the pre-existing channel activity; here, the variability in the RMS trace has decreased significantly, indicating a change in CBSD occupancy behavior, though the duration of its "off" time remains unchanged. By 12.5 min., the CBSD activity has stopped, and the pulse behavior has changed. The pulse repetition interval, at approximately 1 ms, matches the pulse repetition interval "Radar 1," though this is not by itself a conclusive indication of that this specific type of incumbent user is present in the channel.

When combined with other data streams, a PFP analysis like the above represents a powerful tool to add confidence to

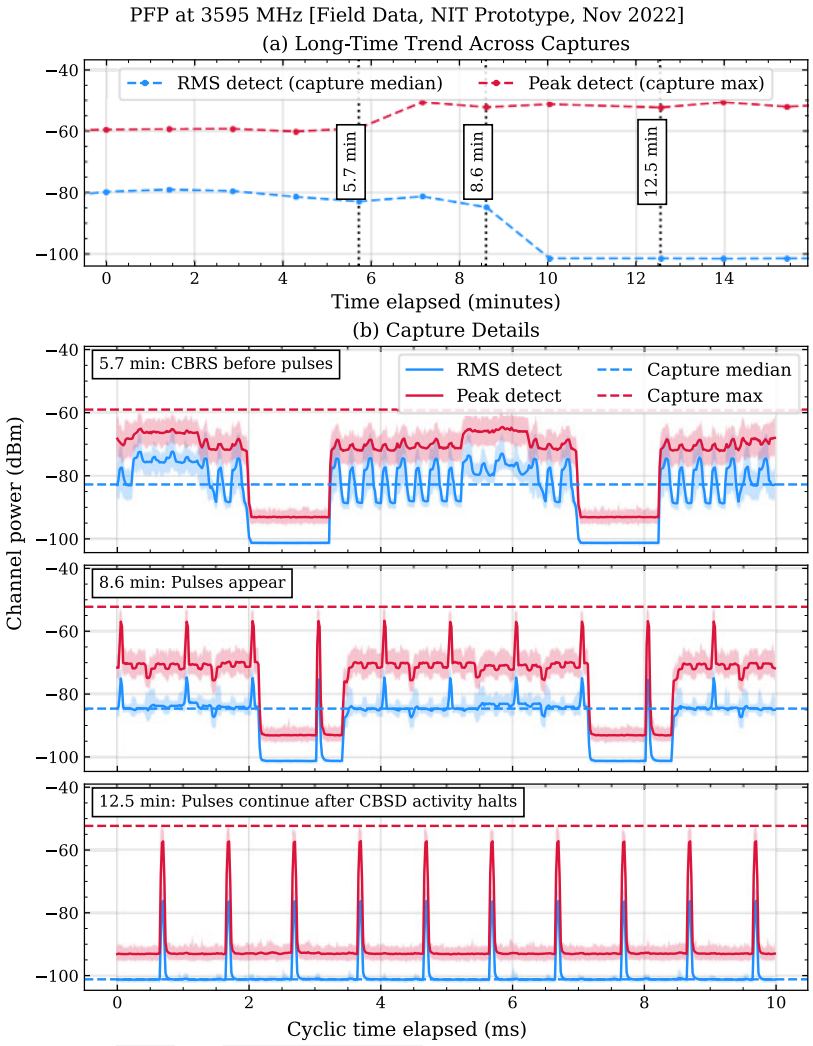

Figure 6.17: Measurements of PFP during a possible Dynamic Protection Area (DPA) activation event. First, time series of capture-by-capture statistics (a) exhibits an abrupt drop in median RMS power. On further investigation, PFP traces in (b) detail CBSD activity before the event ("Capture 1"), followed by the appearance of pulses with 1 ms pulse repetition interval ("Capture 2"), and later the same pulse pattern but without CBSDs. The first capture with pulse activity occurred approximately 3 min before the first capture without CBRS.

conclusions about the CBRS environment. For example, comparisons against logs of SAS channel grants lead CBSDs to respond as expected in terms of physical channel radio channel occupancy. In test cases where "ground truth" data is available that can confirm the presence of incumbents in the sensor's DPA, the PFP can be used to help determine whether power detected by the SEA sensor originates from incumbent or CBSD transmissions.

6.1.3.4 Amplitude Probability Distribution

Description: An APD is the complementary cumulative distribution function (CCDF) of the instantaneous channel power, evaluated empirically. It is evaluated across the full duration of each channel waveform capture. Physically, it is equivalent to the percentage of time in which the capture exceeds a specified power threshold. Plots of the APD provide qualitative insights to the duty cycle of a measured signal. Rare occurrences of high amplitude samples across a full capture may indicate the presence of a strong, briefly pulsed signal, such as pulsed radar. Quantitative analysis

may also leverage the APD to estimate occupancy characteristics in the channel. For example, assuming bursted noise-like transmissions in CBRS wireless data networks, the transmission time-occupancy rate may be estimated based on the deviation between the APD and an assumed underlying distribution, such as idealized Gaussian noise.

Background: The original application for the APD was for the characterization of noise originating in the atmosphere at low frequencies [24]. Because the APD makes no assumption about time structure of the electromagnetic environment, it can be computed independent of wireless networking parameters, unlike periodic frame power (PFP). In conventional graphical representation, the vertical axis is power transformed to decibel units; the horizontal axis is the "percent exceeding ordinate" with a transformation to $Q - Q$ probability space such that idealized thermal noise can be represented as a straight line.

We calculate a discretized APD following Achatz *et al* [25]. The continuous distribution is expressed in terms of the random variable $\mathcal{X}$, a function which assigns a real number $\mathcal{X}(\nu)$ to each element ν of a sample space. The APD is expressed in terms of the probability operator $\mathcal{P}$ and signal amplitude a by Equation 6.9.

$$\mathcal{F}(a) = \mathcal{P}(\mathcal{X}(\nu) > a) \tag{6.9}$$

The discretized APD $\mathcal{F}[n]$ is obtained from the input of N complex baseband samples, IQ$[n]$. First, the amplitudes are computed from the IQ samples, as in Equation 6.10, and sorted in order from smallest to largest:

$$a[n] = \mathrm{sort}\big(\big|\mathrm{IQ}[n]\big|\big). \tag{6.10}$$

Finally, the probabilities are defined based on these sorted amplitudes by Equation 6.11.

$$\mathcal{P}[n] = \mathcal{P}\big(\mathcal{X} > a[n]\big) = 1 - \frac{n}{N} \tag{6.11}$$

This method of obtaining a discretized APD is simple and accurate, yet not practical for continuous real-time sensing. The operation to sort all samples as per Equation 6.10 is computationally costly, which slows down the sweep time across the band. Additionally, the APD output vector has, for both axes, the same dimension as the input vector. This situation is (1) far too large to stream back over the available backhaul solution, at 56 megasamples per 10 MHz channel, and (2) far more detailed along the probability axis than is useful. In a successful effort to optimize IQ post-processing time and output data size, an optimized method was implemented which is based on the definitions of the APD estimate above. In the optimized routine, amplitudes are replaced with amplitude bins, defined based on a parameterized power bin size. The bin edges are defined to span from the minimum amplitude value up to (potentially including but not exceeding) the maximum amplitude value, with uniform spacing by the determined bin size. The output size of the optimized APD therefore depends both on the choice of bin size parameter as well as the dynamic range of the recorded signal, and varies from one capture to the next. Then, the indices are found in the sequence of bin edges at which the unsorted amplitude samples could be inserted while maintaining the sorted ordering. The numbers of occurrences of each index are then counted to give the distribution of amplitude values within in each bin. Finally, these occurrences are cumulatively summed and subtracted from the input length N, before being normalized. This method removes the requirement for a large-input sort operation, and solves both previously mentioned issues. The result is a deliberate trade-off, where resolution is reduced in favor of faster computation and less output data per capture.

Calculation: The APD is computed for each single-channel waveform capture as follows:

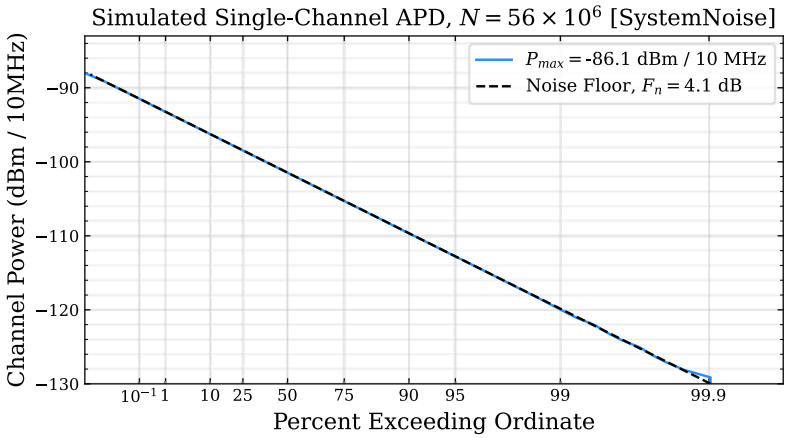

Figure 6.18: Downsampled APD result for simulated system noise in a 10 MHz channel, with simulated channel noise figure $F_n = 4.1$ dB. The predicted APD of the system noise is shown as a black dashed line.

1. Scale IQ samples by the calibrated sensor gain per Equation 6.2.
2. Filter IQ samples using the IIR filter described by Equation 6.3 and Table 6.2.
3. Calculate IQ sample amplitudes by taking the absolute value of each sample, $\left|\mathrm{IQ}[n]\right|$.
4. If any amplitudes are zero-valued, replace them with NaN to avoid floating point arithmetic errors.
5. Generate the sequence of amplitude bin edge values, beginning with the smallest amplitude sample and increasing in uniform steps of half the parameterized power bin size, up to but not exceeding the maximum amplitude value.
6. Using a binary search, locate the indices in the sequence of amplitude bin edges at which the unsorted amplitudes could be inserted while maintaining the sorted bin ordering.
7. Count the number of occurrences of each located index value, representative of the number of amplitude samples contained in each bin.
8. Cumulatively sum the sequence of these counts, and subtract each value from the number of input IQ samples N.
9. Normalize the resulting sequence by dividing each value by the number of input IQ samples N. The results are CCDF probabilities, within the interval $[0, 1]$.
10. Scale the resulting amplitude bin values from dBV amplitudes to power samples in dBm, accounting for system impedance and the power-scaling factor of 2 associated with the down-conversion of the RF signal to baseband.

Examples:

6.1.4 Metadata

Metadata are recorded alongside each edge data sweep in a JSON-formatted text file based on SigMF [15]. The metadata provide the context needed to parse and interpret the raw data, including summary statistics and diagnostic sensor information. Metadata are categorized into global metadata, channelized metadata, and sensor diagnostics.

Global metadata fields contain information pertaining to the entire edge data acquisition, rather than to only a single channel measurement. These are listed in the example JSON file excerpt in Figure 6.21. Not shown in the figure, the global metadata also includes RMS and peak channel power summary statistics, selected details of processing parameters and data collection, and the data component lengths and indices for parsing the binary data file.

Channel metadata fields contain information pertaining to the individual channel measurement. These fields include sensor overload flags, IQ recording timestamps, signal analyzer settings, channel calibration data, and sample index values for parsing the attached data file. The overload flag is triggered on a per-channel basis by the signal analyzer's

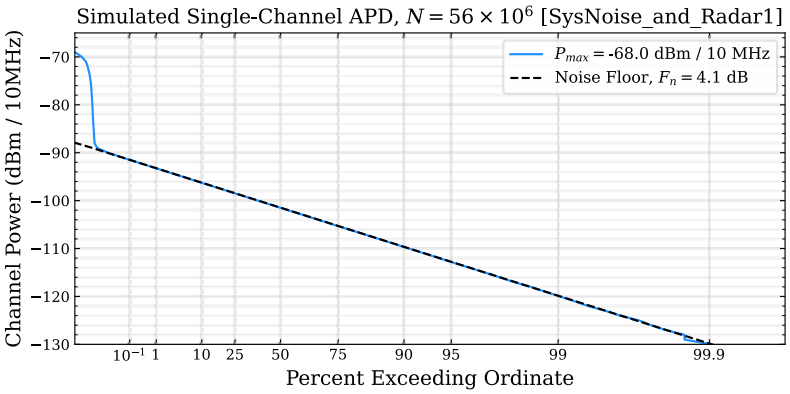

Figure 6.19: Downsampled APD result for simulated system noise and "Radar 1" in a 10 MHz channel, with simulated channel noise figure $F_n = 4.1$ dB. The predicted APD of the system noise is shown as a black dashed line.

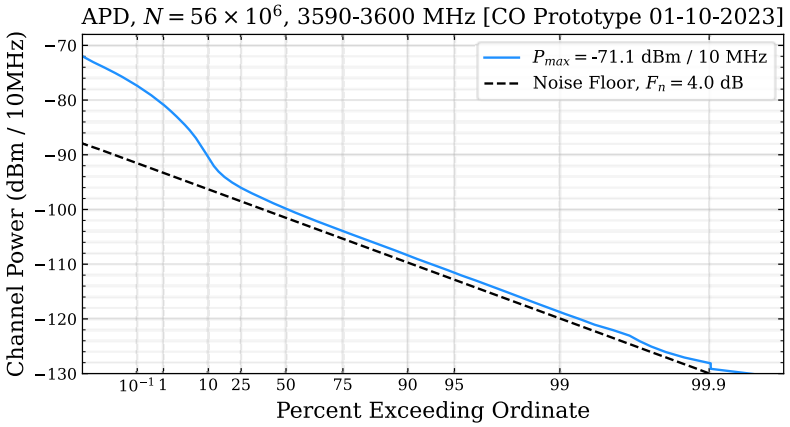

Figure 6.20: Downsampled APD result for real-world signals in the 3590-3600 MHz channel, with the predicted APD of the system noise shown as a black dashed line. The data was recorded on January 10, 2023 by the prototype sensor at Green Mountain Mesa in Colorado.

report of an analog-to-digital converter (ADC) overrange event; see section 4.4 for more details on the conditions that cause over-range. The sample index information is required for aggregation and analysis of the data payloads, since all data is stored in a contiguous binary format.

Diagnostic metadata fields contain information pertaining to the sensor status and operation. These include computer load information and readouts from temperature and humidity probes in various parts of the sensor.

6.1.5 Data Size and Structure

The digital samples representing the results of the IQ processing described above are saved as half-precision (16-bit) floating point binary values. The sample array is converted to binary and compressed using lossless Lempel-Ziv-Markov chain algorithm (LZMA) compression. Using this method, a typical compression ratio of approximately 2.4:1 is achieved, yielding a data file of approximately 70 kB from each measurement of the entire CBRS band. The corresponding data (compressed binary) and metadata (text JSON) files are packaged into a single tar archive. After this file is generated, it is pushed back to a centralized NASCTN data repository over a private LTE network for further cataloging, post-processing, and aggregate analysis. If any error occurs during this file transfer, the data is

```
"global": {
    "core:version": "v0.0.2",
    "core:extensions": {
        "ntia-algorithm": "v1.0.0",
        "ntia-core": "v1.0.0",
        "ntia-environment": "v1.0.0",
        "ntia-location": "v1.0.0",
        "ntia-scos": "v1.0.0",
        "ntia-sensor": "v1.0.0",
        "ntia-nasctn-sea": "v0.3"
    },
    "core:datatype": "rf16",
    "core:sample_rate": 14000000.0,
    "core:num_channels": 15,
    "ntia-scos:task": 1,
    "ntia-scos:schedule": {
        "name": "example_data_product",
        "start": "2023-01-04T23:47:59.000Z",
        "stop": null,
        "interval": 90,
        "priority": 10,
        "id": "example_data_product"
    },
    "ntia-sensor:calibration_datetime": "2023-01-04T23:49:00.311Z",
    "calibration_temperature_degC": 0.60009765625,
    ...
}
```


Figure 6.21: Example of global metadata fields contained in JSON metadata files

saved locally to the sensor flash storage, where it can be later manually retrieved.

6.2 Aggregation

The sensor data streams recorded from each sensor need to be aggregated into a centralized location in order analyze the results across time and location, and to disseminate the data. The detailed approach to this will be developed by NASCTN concurrently with data collection in the field.

Data aggregation will initially consist of a local database hosted on a server accessible for downloading raw data products or via a web interface to be developed at a later time. Periodically, new records received from the streams of sensor data will be parsed and ingested into a database. In this form, the data is available for rapid querying to support subsequent post-processing. In this form, batch data acquisitions can be performed by NASCTN or external staff as needed on data that aggregated across time and sensor site. Key quantities derived from the sensor data can also be added to the database in order to allow for rapid retrieval of pre-calculated measurements.

Chapter 7

Calibration, Verification, Acceptance Tests

7.1 Amplitude Calibration

The sensor system is intended to measure power in the Citizens Broadband Radio Service (CBRS) frequency allocation (3550-3700 MHz). To ensure the power arriving at the input port is accurately reported by the sensor, an amplitude calibration procedure has been designed. The amplitude calibration is a combination of an onboard radiometer-characterized noise diode and a programmable noise probe. The onboard noise diode allows the system to update the system power calibration and provides diagnostic information about the gain and noise figure of the system as deployed. The programmable noise probe functions as a variable noise signal generator with known noise power levels. The programmable noise probe is used to calibrate the sensor at the input port of the preselector without the need of dissembling the sensor. In Figure 7.1, a simplified block diagram of the sensor is shown. The antenna port is labeled as reference plane 1, the output of the on-board calibration noise diode is labeled as reference plane 2. To accurately measure power arriving at the calibration plane labeled C from the plane marked 1, we need to calculate the transducer gain ratio $G_T^{2\to C}/G_T^{1p\to C}$.

7.1.1 Calibration of the On-board Calibration Noise Diode

Both the programmable noise probe and the on-board calibration noise diode are calibrated with noise temperatures, T_p and T_{ND} respectively. Note that the noise diode (ND) has two known states; hot (voltage source on) and cold (ambient, voltage source off). The power measured by the signal analyzer is given by the equations [26]

$$P_p = k_B B \left[G_T^{1p\to C} T_p + M_{C-p} T_s + (M_{C-p} - G_T^{1p\to C}) T_a \right] G \tag{7.1a}$$

$$P_{NDh} = k_B B \left[G_T^{2\to C} T_{NDh} + M_{C-NDh} T_s + (M_{C-NDh} - G_T^{2\to C}) T_a \right] G \tag{7.1b}$$

$$P_{NDc} = k_B B M_{C-NDc} \left(T_s + T_a \right) G \tag{7.1c}$$

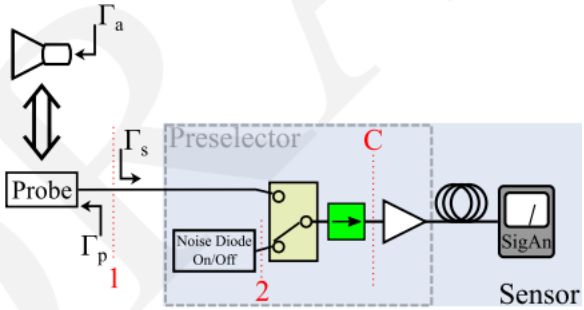

Figure 7.1: The calibration block diagram, identifying the reference plane of antenna input (1), the reference plane of the calibration diode (2) and the calibration plane (C). The probe is inserted into the antenna port with reflection Γ_p, and then is replaced by the antenna in the field with reflection Γ_a, the reflection of the sensor input is Γ_s

Here k_B is the Boltzmann constant. B is the bandwidth set by the signal analyzer or the post-processing digital signal processing (DSP) (usually the frequency bin size). T_s is the equivalent noise temperature of the sensor at plane 'C'. G is the deliverable gain from 'C' to the signal analyzer. The inclusion of the isolator (the green component in Figure 7.1) allows us to approximate the mismatch factors at plane 'C' for dissimilar input terminations $M_{C-p} \approx M_{C-NDh} \approx M_{C-NDc} \equiv M$. As a result, (7.1) can be approximated as

$$P_p = k_B B G_T^{1p\to C} (T_p - T_a) G + k_B B M (T_s + T_a) G \tag{7.2a}$$

$$P_{NDh} = k_B B G_T^{2\to C} (T_{NDh} - T_a) G + k_B B M (T_s + T_a) G \tag{7.2b}$$

$$P_{NDc} = k_B BM\,(T_s + T_a)\,G \tag{7.2c}$$

The three linear equations in (7.2) provide a solution of the transducer gain ratio $\frac{G_T^{2\to C}}{G_T^{1p\to C}}$. We label it as β.

$$\beta = \frac{(P_{NDh} - P_{NDc})(T_p - T_a)}{(P_p - P_{NDc})(T_{NDh} - T_a)} \tag{7.3}$$

In the field, the parameter of interest is T_x, the field radiation temperature measured at the antenna port. When the antenna is connected to the sensor, the transducer gain from '1' to 'C' is $G_T^{1a\to C}$. Similar to (7.2a), the power measured at signal analyzer in the field is

$$P_x = k_B B G_T^{1a\to C}(T_x - T_a)G + k_B BM(T_s + T_a)G \tag{7.4}$$

In combination with the measurements of the programmable noise probe in the field, we obtain:

$$T_x = \frac{G_T^{2\to C}}{G_T^{1a\to C}} \frac{P_x - P_{NDc}}{P_{NDh} - P_{NDc}}(T_{NDh} - T_a) + T_a \tag{7.5a}$$

$$= \beta \frac{G_T^{1p\to C}}{G_T^{1a\to C}} \frac{P_x - P_{NDc}}{P_{NDh} - P_{NDc}}(T_{NDh} - T_a) + T_a \tag{7.5b}$$

Note that the T_a's in (7.3) and the T_a's in (7.5) in practice may deviate substantially. T_a's in (7.3) are measured in the laboratory. β is determined in the lab experiments prior to the sensor deployment. T_a's in (7.5) are frequently updated by temperature monitoring devices in the field. In addition, P_{NDh} and P_{NDc} in (7.5) are calibrations performed in the field.

We determine the ratio $\frac{G_T^{1p\to C}}{G_T^{1a\to C}}$:

From network analysis, the transducer gain relates to the reflection coefficients of the input termination and the receiver. We can express the ratio by [27]

$$\frac{G_T^{1p\to C}}{G_T^{1a\to C}} = \frac{(1 - |\Gamma_p|^2)/|1 - \Gamma_p\Gamma_s|^2}{(1 - |\Gamma_a|^2)/|1 - \Gamma_a\Gamma_s|^2} \tag{7.6}$$

In virtue of (7.6), we represent (7.5) by

$$T_x = \beta \frac{(1 - |\Gamma_p|^2)|1 - \Gamma_a\Gamma_s|^2}{(1 - |\Gamma_a|^2)|1 - \Gamma_p\Gamma_s|^2} \frac{(P_x - P_{NDc})}{(P_{NDh} - P_{NDc})}(T_{NDh} - T_a) + T_a \tag{7.7}$$

7.1.2 Discussion

This treatment demonstrates that the power reading on the sensor is linearly related to the impinging signal temperature. To facilitate understanding, we may combine and regroup Equation 7.7, as a simple linear equation 7.8. We begin by defining the term α equal to the $\alpha = \beta \frac{(1-|\Gamma_p|^2)|1-\Gamma_a\Gamma_s|^2}{(1-|\Gamma_a|^2)|1-\Gamma_p\Gamma_s|^2} \frac{(T_{NDh}-T_a)}{(P_{NDh}-P_{NDc})}$ after algebraic manipulation we are left with

$$P_x = \frac{T_x}{\beta\alpha} + P_{NDc} - \frac{T_a}{\beta\alpha} \tag{7.8}$$

or more succinctly

$$P_x = AT_x + B \tag{7.9}$$

Where A is related to the effective gain and B is a gain scaled constant that represents the ambient noise temperature and the noise temperature of the sensor. The calibration procedure quantifies A and B at a given moment in time at a specific ambient temperature. It is assumed that both of these coefficients can change with time as evidenced by Figure 4.15 (due to ageing and finite system stability) and with temperature (changes to system gain). The onboard calibration noise diode integrated into the sensor allows us to quantify these changes. The more common linear form is

$$P_x = G_{\text{eff}} \Delta F k_B T_x + G_{\text{eff}} \Delta F k_B (T_{ambient} + T_{sensor}) \tag{7.10}$$

Where P_x is the reading of the power sensor, G_{eff} is the effective gain in linear units, ΔF is the bandwidth, k_B is Boltzmann's constant, T_x is the signal noise temperature, $T_{ambient}$ is the noise temperature at the ambient physical temperature, and T_{sensor} is the noise temperature added by the sensor. For simplicity an estimator of G_{eff} was used instead of the full form derived from 7.7. The estimator ignores the difference in impedance matching from the programmable noise diode attached to the input of the sensor and the antenna. In addition, the estimator of G_{eff} is found by using the on-board calibration noise diode installed in the preselector of the sensor. This path is assumed to have a small but fixed difference to the antenna path. For the purposes of this test plan, data collected by the sensor has been adjusted using this estimator and not the full form. The data presented in this test plan thus has higher uncertainties in power than expected in the long term. Production level Sharing Ecosystem Assessment (SEA) sensor will follow the full form calibration procedure, adding to expectations of measurement rigor, and inclusion of terms ignored by the simple estimator.

7.1.2.1 Calibration of Sensor Data

Periodically, the sensor performs a self calibration to estimate G_{eff}, the sensor noise figure, and to update P_{NDh} and P_a. The self calibration procedure first changes the position of the switch in Figure 7.1 from the antenna input port to the onboard noise diode path. Then the noise diode is alternatively energized and deactivated. A recording of in-phase and quadrature (IQ) samples is made at 14 mega samples per second (Msps) for 1 s in each state, then the process is repeated in each 10MHz channel. Each sequence of IQ samples is then filtered using a 10 MHz band-pass filter with a known equivalent noise bandwidth (ENBW). The implementation of this filter is defined by Equation 6.3 and Table 6.2, and further described in Section 6.1.2.3. Then, power is calculated from each recording using the method defined in Equation 6.4 and outlined in subsubsection 6.1.3.1. The ratio of these two powers are then used to calculate the nominal gain and noise figure of the sensor using the Y-Factor method [28] (without the correction for β), as shown in Equations 7.11 through 7.13. These calibration values are stored and reported with each data product measurement. The effective gain is used to scale the recorded IQ transforming them into physical units.

$$y = \frac{P_{\text{on}}}{P_{\text{off}}} \tag{7.11}$$

$$NF = \frac{\text{ENR}}{y - 1} \tag{7.12}$$

$$G_{\text{eff}} = \frac{P_{\text{on}}}{k_B T B (\text{ENR} + \text{NF})} \tag{7.13}$$

Here, P_{on} and P_{off} are the recorded root mean square (RMS) powers in linear units when the noise diode is energized and deactivated, respectively. The noise diode's excess noise ratio (ENR) and noise figure (NF) are computed in these equations in linear units. The noise factor is the noise figure in linear units. In Equation 7.13, k_B is the Boltzmann constant, T is the temperature of the noise diode, and B is the ENBW of the digitally-implemented band-pass filter. Once the calibration is estimated subsequent IQ streams recorded by the spectrum analyzer are scaled by Equation 6.2, using the values calculated above. The last calibration performed is applied until another is executed, for currently deployed sensors this is on an hourly basis. The NASCTN team will monitor for the calibration stability over time to

further adjust the calibration cadence as needed.

7.2 Measurement Uncertainty

There are two key uncertainties that we consider for the sensor system. The first is the uncertainty in the reported power, and the second is the uncertainty of the noise figure estimate. The uncertainty in these values will be refined in future versions of the sensor deployment. The calibration strategy outlined in 7.1, is designed to correct for variations of the reported power for frequency differences, changes due to temperature and the overall path loss difference between the onboard noise diode and the antenna path. Future sensor deployments will have uncertainties calculated using the microwave uncertainty framework (MUF) [29], allowing for the uncertainties in Γ_p, Γ_a, and Γ_s to be accounted for (i.e. the term displayed in 7.2). For the currently deployed prototype sensors an empirical assessment follows.

7.2.1 Reported Power Uncertainty

The uncertainty in reported power is a combination of the uncertainty in the scalar correction of the IQ (see equation 6.2), variability of the sensor in reading a known, fixed power, the difference in mismatch between deployments and any time or temperature variations that are not accounted for by the gain estimator. The sensor measures a known power source (noise diode) to estimate the actual power at the input port for a given reading of the spectrum analyzer. For the three deployed units the on-board calibration noise diode is precisely calibrated in the National Institute of Standards and Technology (NIST) radiometer. The uncertainty in the RMS power emitted from these diodes is extremely small, with ENR uncertainties of approximately 0.05 dB. For other sensor realizations the ENR was derived from the manufacturers measurements with an uncertainty of approximately 0.2 dB. The total sensor power uncertainty is always limited by the uncertainty of the noise diode power. The manufacturer specification for temperature variations of ENR is $0.01\ \frac{dB}{^\circ C}$, and variations in the gain estimator greater than this are assumed to be real changes in system gain and not a change in power emitted from the noise diode. To estimate the variation due to mismatch between the antenna when compared with the variable noise source the term

$$\frac{(1-|\Gamma_p|^2)|1-\Gamma_a\Gamma_s|^2}{(1-|\Gamma_a|^2)|1-\Gamma_p\Gamma_s|^2}$$

is calculated using multiple measurements of the antenna in different laboratory locations. These locations were chosen at random to reflect an potential installation in an uncontrolled environment. In Figure 7.2, the result of this calculation can be seen. This ratio varies from 0.98 to 1.07 indicating a nominal difference of 0.03 from unity.

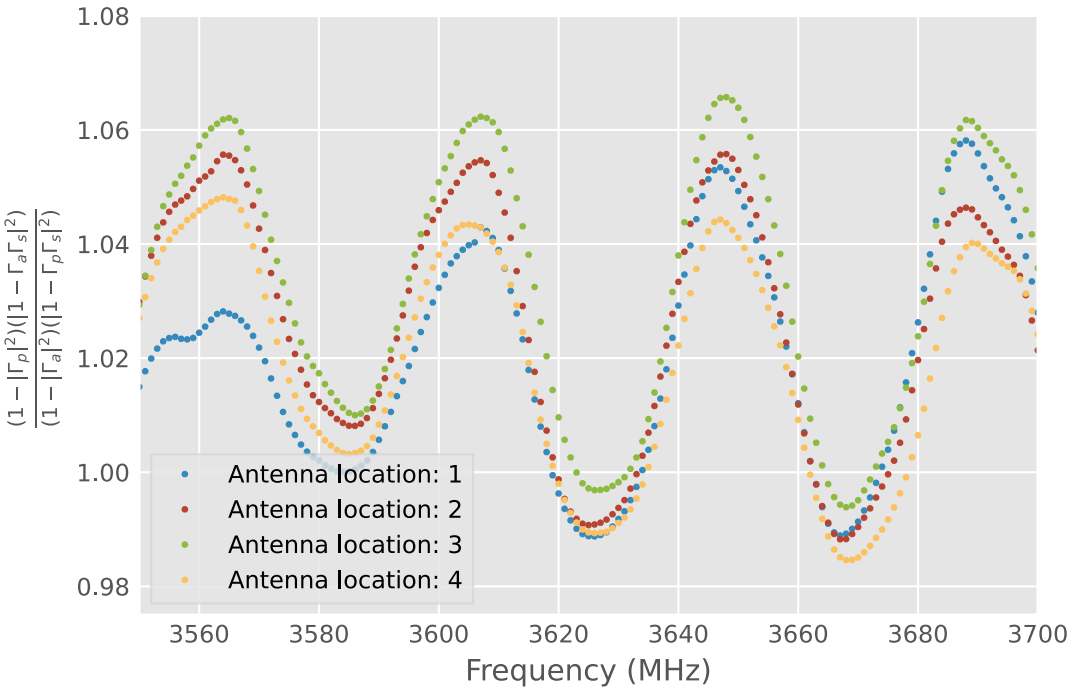

Figure 7.2: Comparison of Antenna Mismatches. The calibrated noise temperature reading requires adjusting for the antenna γ_a and that of the probe γ_p, every installation potential has a slightly different antenna mismatch. To assess the impact of this difference an antenna was measured in four locations.

The temperature variation of the gain can be assessed from field data from the sensor deployed at Hampton University. Data from the time period 10/06/2022 - 12/26/2022 is shown in Figure 7.3 for a single 10 MHz channel, identified as channel 8 with the center frequency of 3625 MHz. It can be seen that the gain in logarithmic units is linearly dependent on the temperature, a fit of this dependence displays a slope of $0.06\ \frac{dB}{^\circ C}$. When the gain values are corrected for this dependence, the residual gain shows a random distribution that we approximate as a Gaussian distribution. The residual gain after correction is shown in Figure 7.4. To better represent the total of gain variations for the sensor that are not corrected for, this process was repeated for all 15 10 MHz channels and aggregated. In Figure 7.5, the results of this process can be seen. The residuals for 597540 measurements shows a standard deviation of 0.15 dB, which leads to an estimate of the expanded uncertainty due to variability of the sensor to measure a known, fixed power and any time and temperature variations unaccounted for as $\pm 0.30\ dB$.

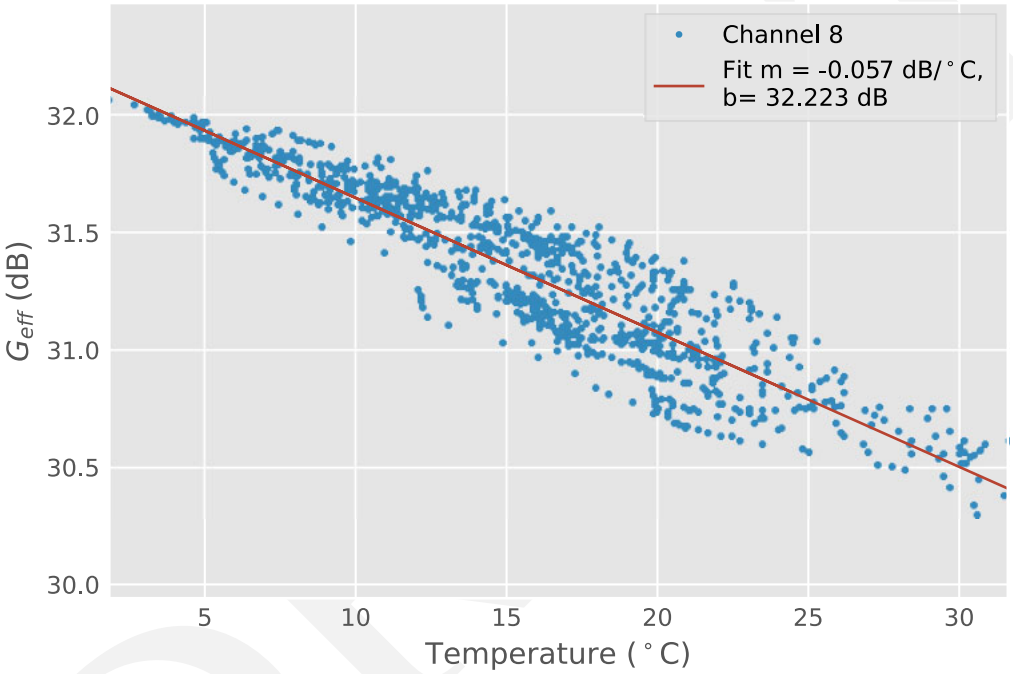

Figure 7.3: Gain (logarithmic units) versus temperature of channel 8 (center f = 3625 MHz) shows a clear linear trend, this is repeated for all channels

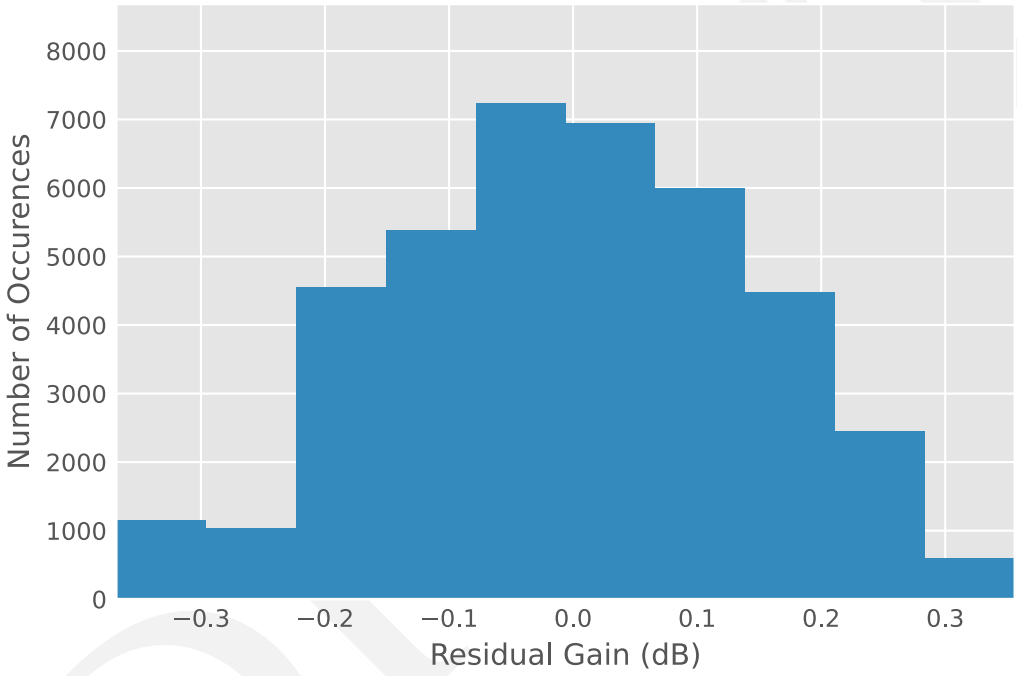

Figure 7.4: Residual Gain After Detrending of Channel 8, Hampton University Sensor.

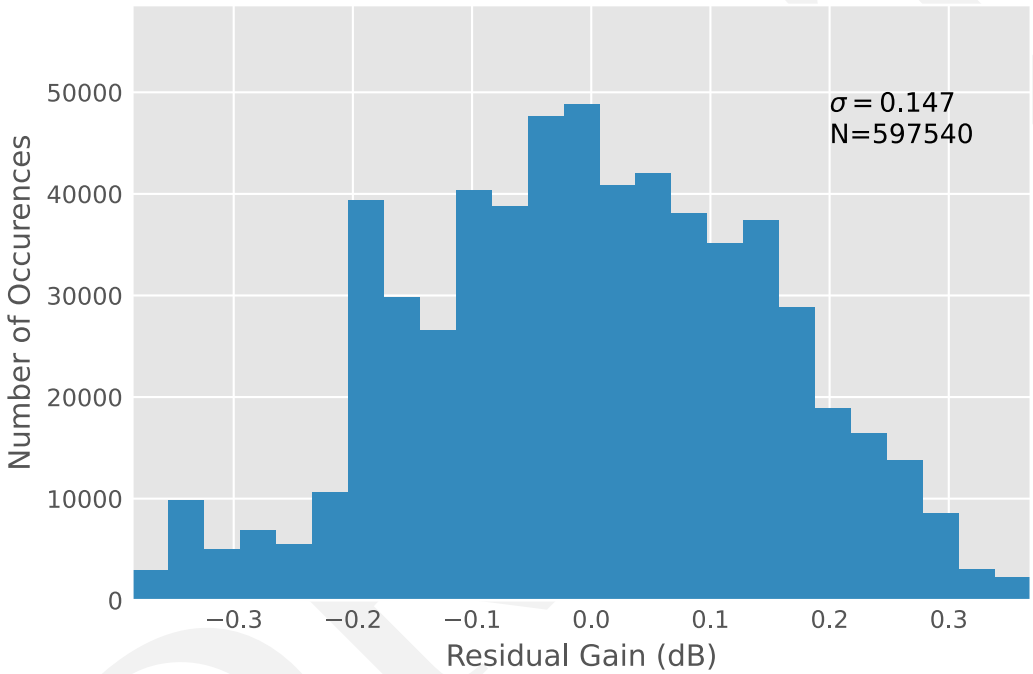

Figure 7.5: Residual Gain after Detrending Temperature Dependence of All Channels, Hampton University Sensor. Histogram of the residual gain after detrending temperature dependence.

7.2.2 Noise Figure Estimate Uncertainty

The noise figure estimate quantifies the amount of random noise added to the signal by the total sensor system, or its decrease in signal to noise ratio (SNR). This value represents the noise floor of the system, and is therefore a valuable metric of performance. The sensor system is designed to be an incoherent (not time synchronized) power detector. The noise figure is estimated by equation 7.12. From this equation, the uncertainty in the onboard noise diode's ENR (0.05 dB) and the uncertainty in the power readings from the sensor when the noise diode is energized and de-energized play a major role in the uncertainty of the noise figure. Any passive loss in the sensor between the antenna and the first low noise amplifier (LNA) will increase the noise figure in portion to the loss. This difference can be seen when comparing 4.32 and 7.6, a 1-2 dB passive loss between the on-board calibration noise diode path and the path leading to the antenna increases the measured noise figure in 4.32 and explains the smaller discrepancy between 4.32 and 4.12. We assume that this difference is fixed and the variation of the noise figure with time and temperature is accurately reflected by the measurement of the on-board calibration noise diode. In Figure 7.6 we show the variation of the estimated noise figure (dB) versus temperature. Similar to G_{eff} the noise figure shows a linear trend with temperature when plotted in logarithmic units, with a slope of $0.013\ \frac{dB}{^\circ C}$ and an intercept of 3.950 for channel 8. In Figure, 7.7, the histogram of the detrended noise figures can be seen. The other channels show similar slopes as channel 8. After detrending all of the channels for their respective slope and intercept, and aggregating the results, a combined standard deviation of 0.02 dB is observed (see Figure 7.8).

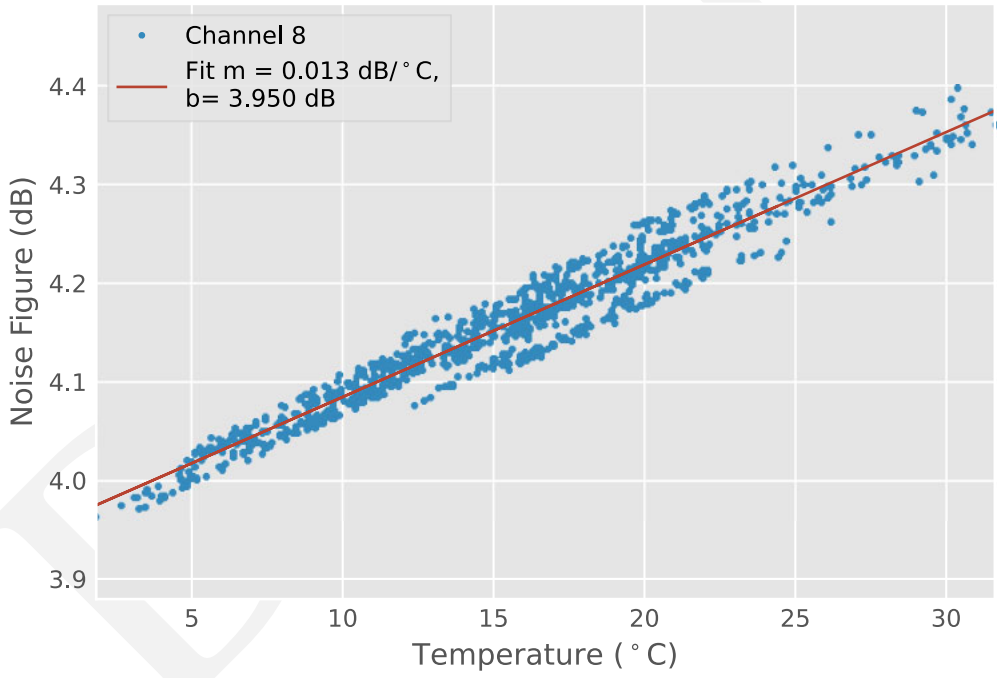

Figure 7.6: Temperature Dependence of Noise Figure channel 8, Hampton University Sensor.

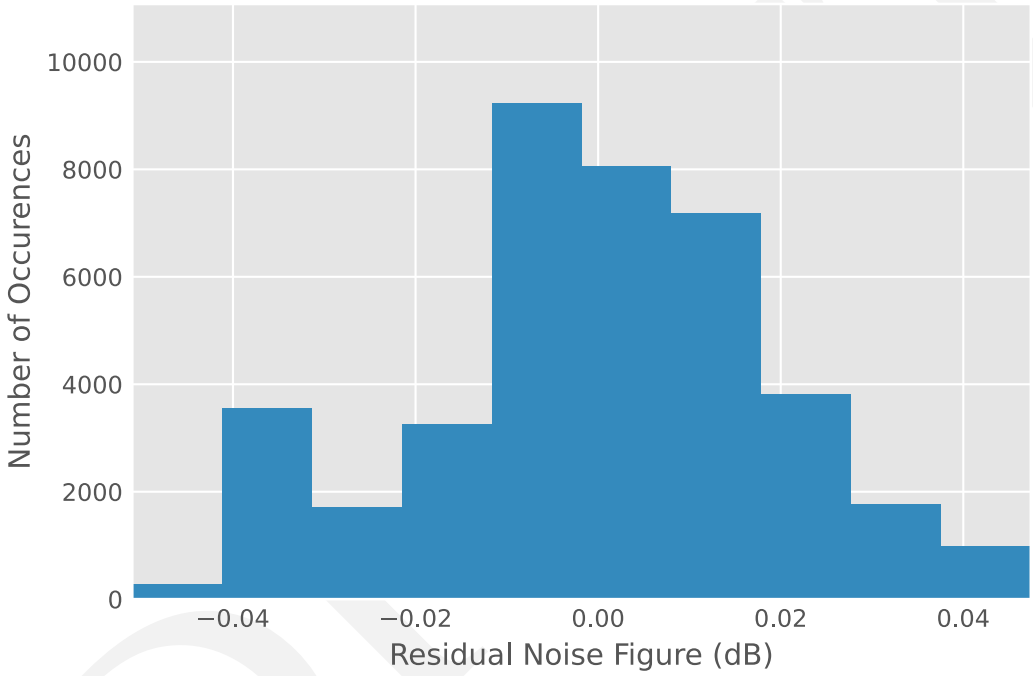

Figure 7.7: Residual Noise Figure after Detrending Temperature Dependence channel 8, Hampton University Sensor.

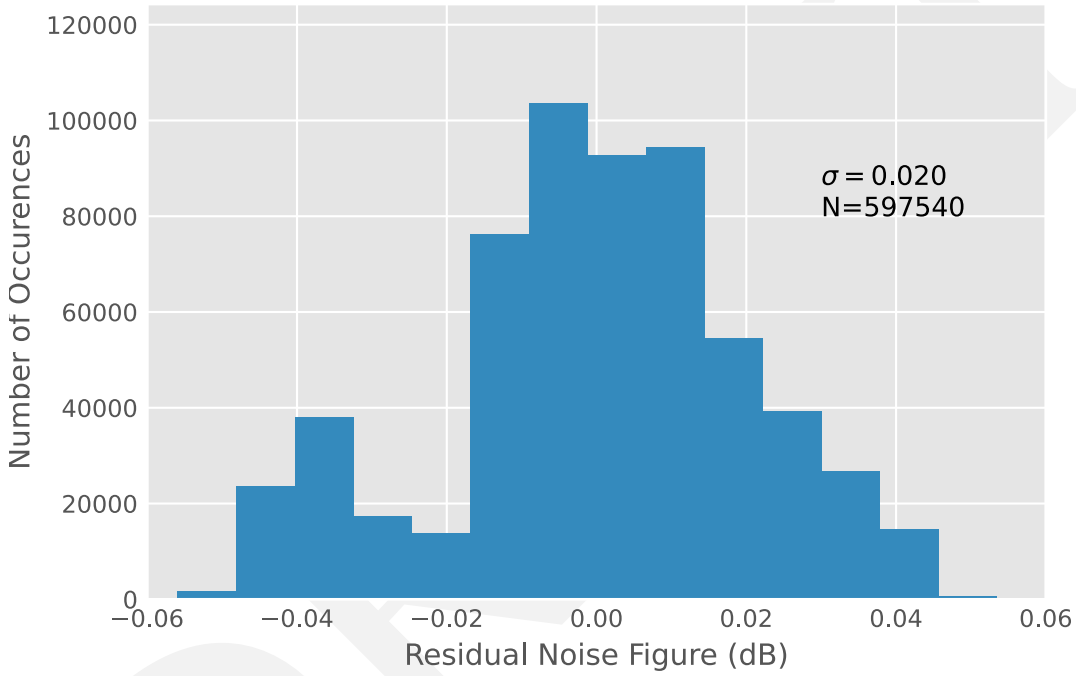

Figure 7.8: Residual Noise Figure after Detrending Temperature Dependence for All Channels, Hampton University Sensor. This histogram of the residual noise figure after detrending temperature dependence shows a standard deviation of 0.02 dB

7.3 Sensor Verification and Acceptance Tests

Once the sensor was constructed and identified as operational, a series of tests for acceptance were administered. These tests were chosen to be system level verification and not necessarily component level tests. The verification was three-fold:

1. RF characterization
2. Software verification
3. System Power and communication verification

Only once these three sections were completed, were the sensors considered fully deployed at the field locations.

7.3.1 RF Characterization

For RF characterization, a calibration probe, consisting of an amplified noise diode with a programmable attenuator was devised. This essentially operated as a noise diode with a settable ENR. A bench top sensor system used for initial prototyping was used to determine the attenuation setting versus ENR for the calibration probe.

Once the probe was calibrated, the completed sensors were measured with the calibration probe with a fixed ENR/attenuation from the antenna terminal in contrast to the on-board noise diode. This was done prior to deployment (see Figure 7.11) and at several stages during installation to ensure nothing was damaged during transit or assembly.

Finally, the S_{11} of the sensor and installation cable was taken before deployment. Then, once installed in the field, the S_{11} of the antenna was measured just before connection to the final system to ensure the antenna's environmental radiation pattern was accounted for.

7.3.2 Software Verification

In the lab, prior to deployment, the spectrum-characterization and occupancy sensing (SCOS) software was used to run a number of actions to get a distribution of trends for each sensor. The sensor was connected end-to-end with exception of the antenna, which is either replaced with the probe or left unterminated for calibrations and the software was controlled locally through a breakout Ethernet connection.

First, the calibration action was run ten (10) times and the resulting compiled calibration files are recovered from the sensor. This was used to determine an expected spread of calibration values. The resultant calibration trend was used to compare with post deployment calibrations to ensure the functionality up to the preselector had remained intact. An onboard calibration was taken at several points during the installation process up to the point when the preselector was mounted. The field data was verified to be similar to the lab-taken data, see Figure 7.9.

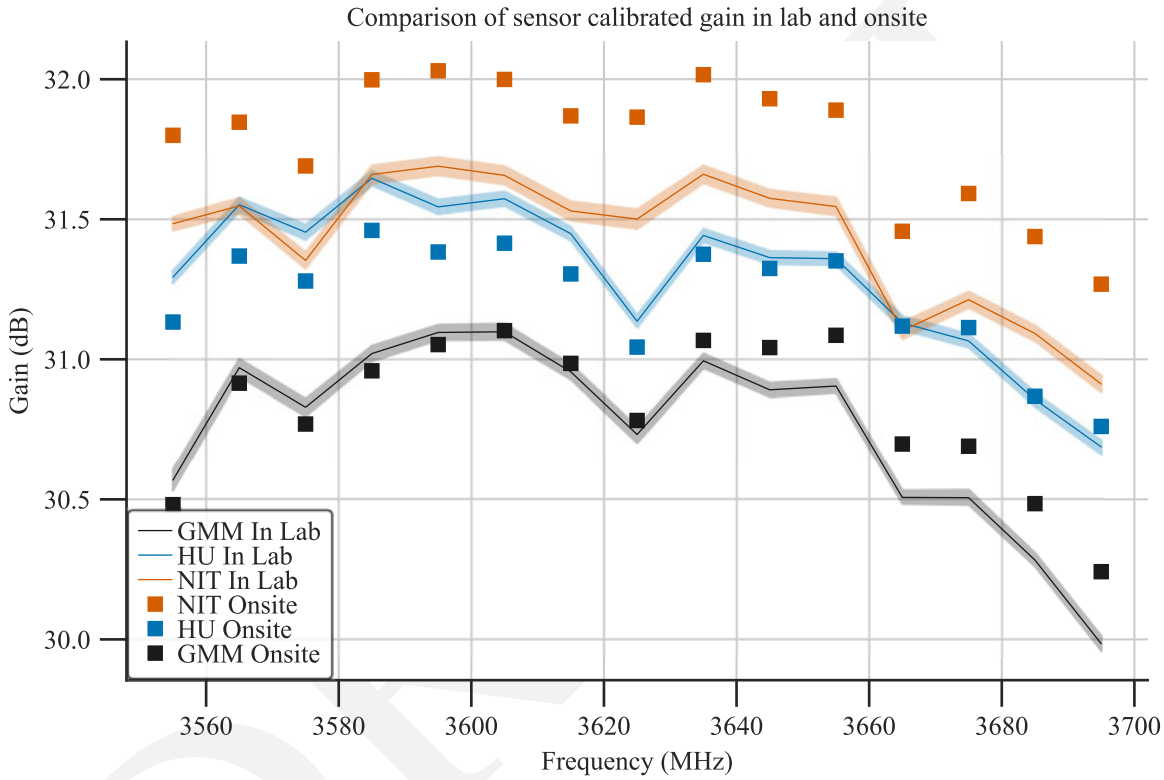

Figure 7.9: Gain Calibration in the Laboratory and as Deployed. Each unit had 10 calibration actions performed in rapid succession in the laboratory, the solid line is the median of those calibration while the shaded region is 2σ. The solid squares show the response in as installed for 1 calibration action.

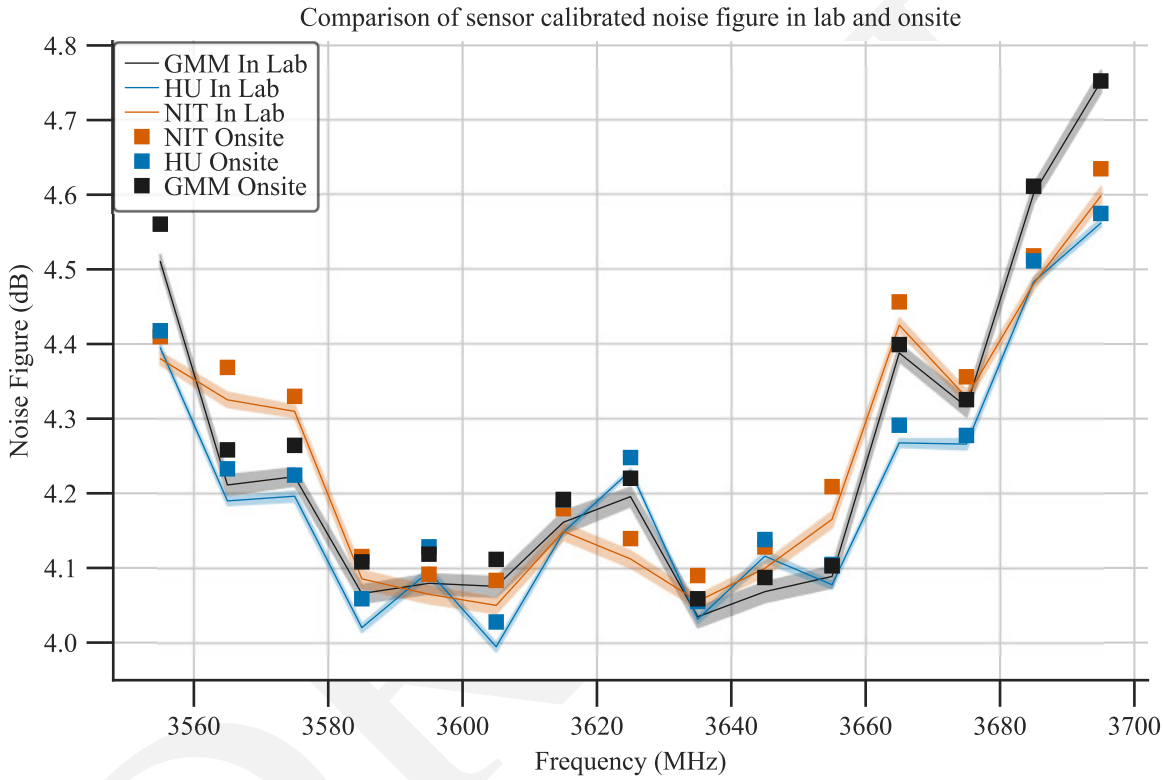

Figure 7.10: Noise Figure Calibration in the Laboratory and as Deployed. Each unit had 10 calibration actions performed in rapid succession, the solid line is the median of those calibration while the shaded region is 2σ. The solid squares show the response in as installed for 1 calibration action.

Similarly, with the above-mentioned calibration probe attached with a fixed attenuation, the data product is taken ten (10) times to get a distribution for the data product as well. Once the sensor is installed, the calibration probe is attached at the antenna port and the data product is taken. It is finally verified that the field-measured data is within the expected range from the lab-taken data to ensure a successful installation.

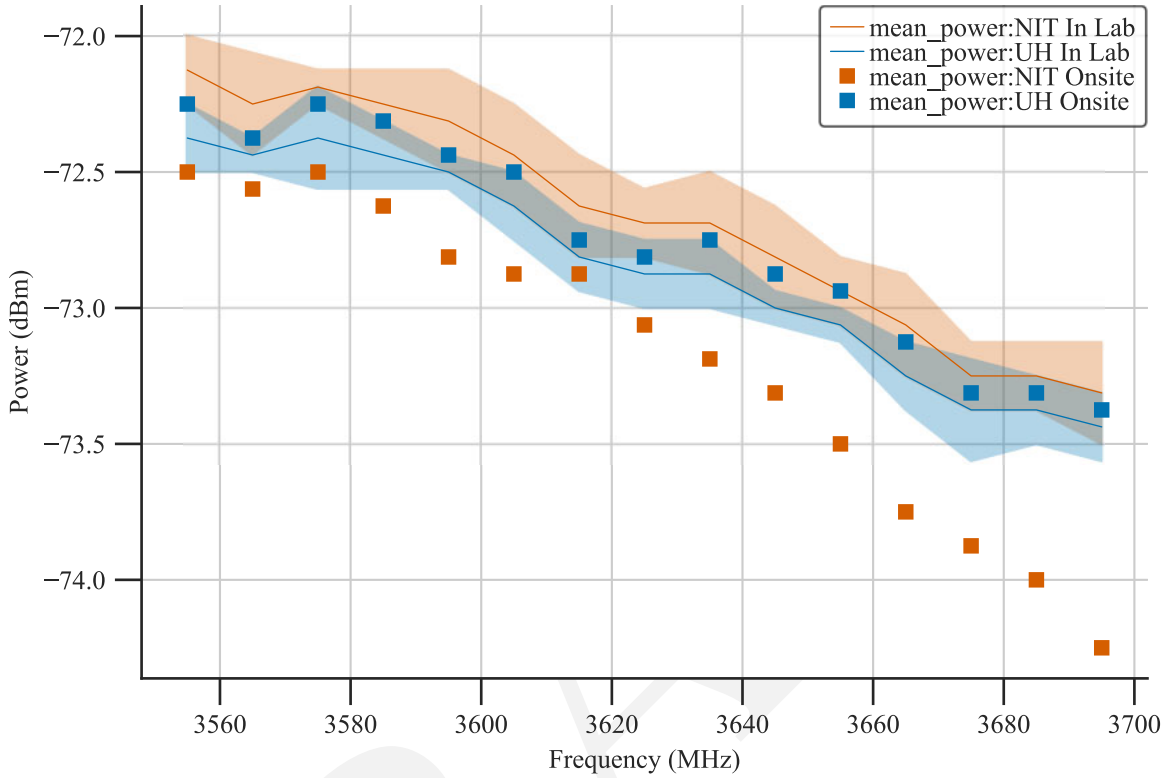

Figure 7.11: Measurement of a programmable noise source inside of the laboratory and at the installation sites.The solid line is the median of ten (10) measurements performed in a laboratory environment, with the shaded region representing the expanded uncertainty (2σ). The squares were measured post installation in an uncontrolled environment.

7.3.3 System Electrical Power and Communication Verification

In the laboratory before installation, if the system presented results consistent with the previous characterizations, the system electrical power was considered fully functioning. However, if there was an issue, there are diagnostic adaptors for each of the connectors and test points throughout the system to test for electrical power problems. Even in full working order, it was still verified that all web-relays could be cycled (with audible feedback if applicable). A final check of the power system was performed to ensure that the the AC web-switch could be cycled and that the controlling computer powered back up once power was restored.

Finally, the data backhaul connection was tested by connecting to it through the local network. It was ensured that each of the 3 switches could be remotely controlled, the router could be accessed, and that the controller computer could be remotely controlled. This was repeated in the field and the systems were registered through SCOS manager to ensure there was a consistent data backhaul for each sensor.

Chapter 8

Metrology Sensor Demonstration

8.1 Objectives

The National Advanced Spectrum and Communications Test Network (NASCTN) team has deployed three prototype Rev 3 sensors in order to demonstrate the efficacy of the approach and systems design described in this test plan. One prototype sensor was deployed locally in Boulder, Colorado in early October 2022. Two additional prototype sensors were deployed in the Hampton Roads region of coastal Virginia in late October 2022. Deployment of these prototype sensors has provided real-world data payloads for analysis. This has also served as a demonstration of the sensor installation process, and as a test case for remote data backhaul links. The data collected has been used to inform iterations on sensor hardware, software design, edge compute products, and system uncertainty characterization. In all cases, sensor installation went smoothly and all three sensors have been able to provide usable data payloads. Previously unforeseen error modes were identified through long-term sensor operation, and have since been addressed through software updates pushed to the sensors. To identify locations for deployment of prototype sensors, site surveys were carried out in advance of sensor installation.

8.2 Site Surveys

8.2.1 Background

Site surveys were performed in order to identify and evaluate potential sensor deployment locations for the NASCTN Sharing Ecosystem Assessment (SEA) project. These surveys aimed to identify infrastructure supporting sensor requirements, such as access to power and antenna mounting points, as well as to provide initial assessments of the radio frequency (RF) environments around potential deployment sites. A portable measurement system was designed for the purpose of conducting RF spectrum measurements in the Citizens Broadband Radio Service (CBRS) band during site surveys. This system had a nominal front end gain of 25 dB and a nominal noise figure of 4.5 dB in the CBRS band. The surveys also provided opportunities to interface with area managers to discuss the required processes and logistics for sensor deployment. Using the results from these surveys in corroboration with modeling-informed analysis, NASCTN is to select suitable sites for sensor deployment in order to meet SEA test goals.

8.2.2 Locations Visited

Marine Corps Base Camp Pendleton was selected for surveying due to its status as an always-on ground-based Dynamic Protection Area (DPA) (GB-DPA). NASCTN team members from National Institute of Standards and Technology (NIST), the Institute for Telecommunication Sciences (ITS), and The MITRE Corporation (MITRE) arrived at Marine Corps Base Camp Pendleton on March 16, 2022, and conducted site surveys over the following three days. In that time period, 22 sites were visited and surveyed for potential deployment. Figure 8.1 shows a map of the surveyed sites.

Various locations in coastal Virginia and North Carolina were visited, spanning approximately 85 miles North-to-South from Cape Charles to Kill Devil Hills. The locations were selected due to their proximity to the East 1 E-DPA The majority of surveyed locations were concentrated in the Hampton Roads metropolitan area. NASCTN team

Figure 8.1: Sites surveyed at Marine Corps Base Camp Pendleton from March 16 to March 19, 2022. The Camp Pendleton GB-DPA is shown outlined in light blue, and the nearby West 14 Environmental Sensing Capability (ESC)-monitored DPA (E-DPA) is outlined in red. All markers indicate surveyed sites, while markers with star icons embedded indicate surveyed sites where RF signals were recorded in the CBRS band above the sensor noise floor.

members from NIST, ITS, and MITRE arrived in the area on March 28, 2022, and conducted site surveys over the following four days. In that time period, eight sites were visited and surveyed for potential deployment. Visited sites included locations operated by the United States Navy (USN), the National Park Service (NPS), the United States Fish and Wildlife Service (USFWS), the National Aeronautics and Space Administration (NASA), the Virginia Marine Resources Commission, and the Virginia Port Authority (VPA). Due to only observing CBRS signals at one location, another potential deployment was identified following the March 2022 surveys in Virginia and North Carolina. A building belonging to Hampton University (HU) was identified as a potential prototype deployment site, and surveyed by a NASCTN team member from NASA Langley Research Center (LaRC) on September 13, 2022. The survey for this location made use of a slightly modified, but comparably performant, measurement system due to local equipment availability. Signals were observed in the CBRS band from this location. Figure 8.2 shows a map of the surveyed sites, both in March and September 2022. In the figure, the cluster of 3 overlapping markers are various sites at the NASA Langley Research Center.

8.2.3 Survey Methods

Both qualitative and quantitative analyses were performed to evaluate each visited site for sensor deployment feasibility. These were driven by expected requirements for a sensor deployment, which include access to alternating current (AC) power, strong long-term evolution (LTE) connection for data backhaul, and proper physical infrastructure to support the antenna and sensor. Pictures were taken at each site, and available facilities and site access details

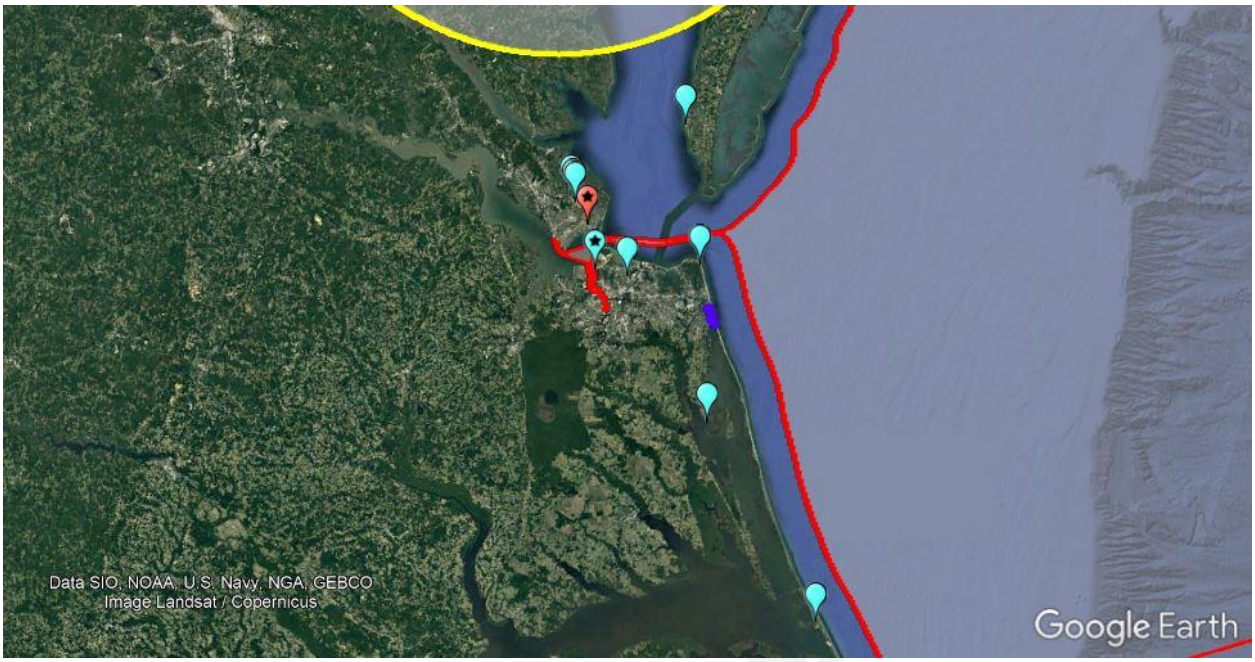

Figure 8.2: Sites surveyed near the East 1 E-DPA from March 28 to March 31, 2022. The East 1 E-DPA is shown outlined in red, the Dam Neck GB-DPA is outlined in light blue, and the St. Inigoes CBRS exclusion zone is outlined in yellow. All markers indicate surveyed sites, while markers with star icons embedded indicate surveyed sites where RF signals were recorded in the CBRS band above the sensor noise floor. Blue markers indicate sites surveyed in March 2022; the red marker indicates the site surveyed in September 2022.

were documented. Cellular uplink and downlink speed tests were performed to evaluate LTE backhaul potential.

8.2.4 RF Survey Results

The primary purpose of the RF surveys at each site was to identify the presence of signals in the CBRS band. The measurement system was calibrated using the Y-Factor method at the start of each day when survey measurements were performed. As shown in Figure 8.1, eleven of the surveyed sites at Camp Pendleton indicated the presence of at least one signal of any type with power above the noise floor of the survey measurement system. In most cases, these signals were low-level, and in some cases, the signals appeared to be incumbent radar systems. As shown in Figure 8.2, surveys recorded signals in the CBRS band at only one site on the East Coast. At this site, the signals appeared to wireless broadband cellular downlink signals, populating the entire band with power levels substantially above the sensor noise floor. Overall, the surveyed sites at which stronger CBRS signals were recorded were considered more preferable for sensor deployments, compared to sites in which no signals were recorded.

8.3 CBRS-ecosystem observations

Prototype Rev3 sensors were deployed in three locations: one in Green Mountain Mesa (GMM) in Boulder, CO, and two in the Norfolk/Hampton, VA area at the Norfolk International Terminal (NIT) tower and HU building. These sensors represent live deployments and offer the opportunity to exercise the sensors in live environments, e.g. formulating data products on the sensors, ensuring a reliable backhaul connection, and failure mode analysis.

When deploying, the signal processing unit (SPU) was placed in an indoor, sheltered location and cabling was run to the preselector installed outdoors near the antenna mount to minimize the cable loss to the antenna. The installation parameters were recorded (Table 8.1) where the heights were recorded as both above ground level (AGL) above mean

Table 8.1: Deployment parameters for the three prototype sensors.

Parameter	GMM Sensor	NIT Sensor	HU Sensor
Latitude	39.99180°	36.91552°	37.02578°
Longitude	-105.2745°	-76.32245°	-76.34120°
Azimuth	12° ± 5°	N/A	N/A
Elevation	0° ± 1°	0° ± 3°	0° ± 3°
AGL Estimate	4.09 m ± 0.05 m	45 m	65.5 m
AMSL (WGS-84)	1783.4 m ± 1.5 m	48.5 m ± 1.5 m	70.4 m ± 1.5 m
Antenna	KP-3SX4-90	HG3509U-PRO	HG3509U-PRO

sea level (AMSL) (World Geodetic System 1984 (WGS-84)).

As the systems were deployed, calibration actions were run at several points and compared to lab calibrations to ensure that the system did not have any faults as it was being assembled and mounted. Once the installation was complete, the calibration probe described in subsection 7.3.1 was used to perform a final measurement at the antenna cable. This measurement was then compared to the lab results to ensure the system displayed similar performance. Finally, the S_{11} of the as-deployed antenna was taken so the impedance difference between the calibration probe and the antenna could be accounted for.

Once the systems were fully deployed, the backhaul connectivity was tested. First, the signal strength indicator was used to position the cellular router and its antennas in the case of the VA sensors. Then, coming in through the laboratory network, it was ensured that each component (computer, preselector, webswitch, etc.) was accessible remotely. A calibration action was then run followed by a data product action to ensure data was remotely accessible. The results of the data product action were also used to compare with prior site survey results for the location. Finally, the sensors were registered in spectrum-characterization and occupancy sensing (SCOS) manager to enable the tasking of calibration actions on an hourly basis and data product actions with data retrieval.

8.3.1 Green Mountain Mesa, Boulder, CO

The first prototype sensor was deployed locally at GMM near the Department of Commerce (DoC) Boulder laboratories campus. This nearby location offered an easily managed sensor where repair or revision turnaround was relatively short and failure modes could be identified and rectified quickly. This allowed it to be used a development sensor, pushing software updates or revising hardware and could be fixed in a relatively short manner if a problem occurred. In addition, this sensor retained a hardline backhaul to the network. This removed any data or bandwidth limitations from the sensor, allowing much more frequent software updates or the capture of raw-form data for testing data products.

The deployment of this prototype sensor offered the ability to exercise the installation plan, including ensuring all necessary equipment was included. Throughout the installation, failure modes and solutions were recorded, e.g. ensuring the RF tray has power after a power cycle. It was also ensured that each installation hurdle had a solution that was remotely configurable to ensure that future deployments would not run into prolonged periods of down time due to these installation challenges.

8.3.2 East 1 E-DPA

Following the surveys near the East 1 E-DPA, prototype sensors were also deployed to two locations in the Norfolk and Hampton, VA areas. The two sensors were installed in late October, 2022. The first was installed atop a tower operated by the Virginia Port Authority within the NIT. The second was installed atop a building operated by HU. The two sites are shown in Figure 8.2 as the blue and red markers with embedded star icons, respectively. Both of

Figure 8.3: Image of deployed prototype Rev 2 and 3 sensors at the GMM site in Boulder, CO. The close proximity to Boulder labs makes this an ideal location to test prototypes before fully remote deployments.

these locations offered high-elevation vantage points above the surrounding area. Site surveys indicated strong CBRS occupancy from both locations. Deployment of these prototype sensors included the full system, including remote setup and operation as well as the cellular backhaul to the central server.

When first installed, the sensors exhibited approximately 90% data success rate. Of the data product and calibration actions, 90% succeeded and returned valid data. After analysis of failures, incremental updates were devised. These were first tested on the GMM deployment, then were rolled out to the VA sensors once stable. After the most recent software updates, NASCTN has not encountered additional failure modes over the most recent three week period.

These sensors offer the opportunity to take preliminary data incorporating both CBRS and the incumbent. The proximity to the Norfolk harbor and East 1 DPAs has the potential for observing incumbent signals as well as the dynamics of the CBRS ecosystem, such as in Figure 6.17. While the focus in this work is on GB-DPAs, the specific Camp Pendleton DPA is in close proximity to the West 14 DPA and will likely observe the dynamics of the CBRS ecosystem in addition to the protections allotted to GB-DPAs.

(a) Sensor staged for deployment (b) Preselector deployed at NIT (c) Deployed sensor at HU

Figure 8.4: Pictures of the prototype sensors deployed in Virginia near the East 1 DPA.

Chapter 9

Conclusions

The goal of this National Advanced Spectrum and Communications Test Network (NASCTN) test plan is multi-fold. To that end, some important elements from the test plan include the following:

- The test plan, though specific to the ground-based Dynamic Protection Area (DPA) (GB-DPA) at Camp Pendleton, encompasses a uniform sensor architecture that is anticipated to be deployed across the Sharing Ecosystem Assessment (SEA) project task elements
- This test plan document serves as a reference for upcoming NASCTN SEA test plans.
- We anticipate, pending adjudicated commentary, the following to stay largely the same across test plans:
  1. Sensor RF system architecture
  2. spectrum-characterization and occupancy sensing (SCOS) based sensor management
  3. site survey methodology
  4. blueprint for sensor system validation, acceptance testing and calibration
- NASCTN is still investigating and developing towards the ingestion and validation of data at scale. The edge compute data products highlighted in this document are considered pre-decisional draft data.
- Modeling and simulation heavily influenced by the Wireless Innovation Forum (WInnForum) reference implementation, which is based on the Irregular Terrain Model (ITM). This does not constitute an endorsement by NASCTN of any particular propagation model but exemplifies a methodology towards incorporating modeling and simulation into the data driven insights needed to quantify effects in the Citizens Broadband Radio Service (CBRS) ecosystem.

Bibliography

[1] LiConn Inc, "Wide Band Low Noise Amplifier," 2014. [Online]. Available: https://www.liconn.com/products/LNA02004000A.pdf (Last accessed: 6 February 2023).

[2] National Telecommunications and Information Administration, "Dynamic Protection Areas." [Online]. Available: https://www.ntia.doc.gov/fcc-filing/2015/ntia-letter-fcc-commercial-operations-3550-3650-mhz-band (Last accessed: 13 April 2021).

[3] CBRS WInnForum Standards, "Signaling Protocols and Procedures for Citizens Broadband Radio Service (CBRS): Spectrum Access System (SAS) - Citizens Broadband Radio Service Device (CBSD) Interface Technical Specification, Document WINNF-TS-0016, Version V1.2.7," 2022. [Online]. Available: https://winnf.memberclicks.net/assets/CBRS/WINNF-TS-0016.pdf (Last accessed: 16 August 2022).

[4] CBRS WInnForum Standards, "Requirements for Commercial Operation in the U.S. 3550–3700 MHz Citizens Broadband Radio Service Band, Document WINNF-TS-0112, Version V1.9.0," 2020. [Online]. Available: https://cbrs.wirelessinnovation.org/release-1-standards-specifications (Last accessed: 07 September 2022).

[5] National Telecommunications and Information Administration, Institute for Telecommunication Sciences, "Irregular Terrain Model, Version 1.2.2." [Online]. Available: https://github.com/NTIA/itm-longley-rice (Last accessed: 05 October 2017).

[6] Wireless Innovation Forum, "SAS Testing and Interoperability Repository." [Online]. Available: https://github.com/Wireless-Innovation-Forum/Spectrum-Access-System (Last accessed: 18 April 2022).

[7] Wireless Innovation Forum, "Common WInnForum Data and Libs." [Online]. Available: https://github.com/Wireless-Innovation-Forum/Common-Data (Last accessed: 13 April 2021).

[8] Federal Communications Commission, "Part 96 - Citizens Broadband Radio Service," 2015. [Online]. Available: https://www.ecfr.gov/current/title-47/chapter-I/subchapter-D/part-96 (Last accessed: 08 February 2023).

[9] E. Drocella, J. Richards, R. Sole, F. Najmy, A. Lundy, and P. McKenna, "3.5 GHz Exclusion Zone Analyses and Methodology," National Telecommunications and Information Administration, Tech. Rep. TR-15-517, 2015. [Online]. Available: https://www.ntia.gov/report/2015/35-ghz-exclusion-zone-analyses-and-methodology

[10] L-Comm, "HyperLink Wireless 3.5 GHz 9 dBi Professional Omnidirectional Antenna." [Online]. Available: https://www.l-com.com/Images/Downloadables/Datasheets/ds_HG3509U-PRO.pdf (Last accessed: 6 February 2023).

[11] Tektronix, "RSA500A," 2023. [Online]. Available: https://download.tek.com/datasheet/RSA500A-Datasheet-final-19Jan2023.pdf (Last accessed: 6 February 2023).

[12] M. Cotton, L. Vu, B. Eales, and A. Hicks, "3.45–3.65 GHz Spectrum Occupancy from Long-Term Measurements in 2018 and 2019 at Four Coastal Sites," National Telecommunications and Information Administration, Tech. Rep. TR-20-548, 2020. [Online]. Available: https://its.ntia.gov/umbraco/surface/download/publication?reportNumber=TR-20-548.pdf

[13] "IEEE Standard for Spectrum Characterization and Occupancy Sensing," IEEE Std 802.15.22.3-2020, pp. 1–58, 2020, (Last accessed: 08 February 2023).

[14] Defense Information Systems Agency. Canonical Ubuntu 20.04 LTS STIG for Ansible - Ver 1, Rel 6. [Online]. Available: https://dl.dod.cyber.mil/wp-content/uploads/stigs/zip/U_CAN_Ubuntu_20-04_LTS_V1R6_STIG_Ansible.zip

[15] The GNU Radio Foundation, Inc. (2022) The signal metadata format (sigmf), v1.0.0. (Last accessed: 08 February 2023). [Online]. Available: https://github.com/gnuradio/SigMF

[16] R. G. Lyons, Understanding Digital Signal Processing, 3rd ed. Upper Saddle River, NJ: Prentice Hall, 2011.

[17] J. W. Cooley and J. W. Tukey, "An Algorithm for the Machine Calculation of Complex Fourier Series," Mathematics of Computation, vol. 19, no. 90, pp. 267–301, April 1965.

[18] R. Dalke, "Radio Spectrum Estimates Using Windowed Data and the Discrete Fourier Transform," National Telecommunications and Information Administration, Tech. Rep. TR-10-470, 2010, (Last accessed: 08 February 2023). [Online]. Available: https://its.ntia.gov/umbraco/surface/download/publication?reportNumber=10-470[2].pdf

[19] D. Kuester, X. Lu et al., "Radio spectrum occupancy measurements amid covid-19 telework and telehealth," 2022-10-14 2022. [Online]. Available: https://tsapps.nist.gov/publication/get_pdf.cfm?pub_id=934697

[20] "5G NR Physical Channels and Modulation, Technical Specification 38.211 (Release 15)," 3rd-Generation Partnership Project, Tech. Rep., 2018.

[21] "LTE E-UTRA Physical Channels and Modulation, Technical Specification 36.211 (Release 15)," 3rd-Generation Partnership Project, Tech. Rep., 2018.

[22] "IEEE Standard for Air Interface for Broadband Wireless Access Systems," IEEE Std 802.16-2017 (Revision of IEEE Std 802.16-2012), pp. 1–2726, 2018.

[23] F. H. Sanders, J. E. Carroll, G. A. Sanders, R. L. Sole, J. S. Devereux, and E. F. Drocella, "Procedures for Laboratory Testing of Environmental Sensing Capability Sensor Devices (Tech. Mem. 18-527)," National Telecommunications and Information Administration, Tech. Rep., Nov. 2017. [Online]. Available: https://its.ntia.gov/umbraco/surface/download/publication?reportNumber=TM-18-527.pdf

[24] R. S. Hoff and R. C. Johnson, "A Statistical Approach to the Measurement of Atmospheric Noise," Proceedings of the IRE, vol. 40, pp. 185–187, 1952.

[25] R. J. Achatz, M. G. Cotton, and R. A. Dalke, "Estimating and Graphing the Amplitude Probability Distribution Function of Complex-Baseband Signals (IEEE P802.15-04-0428-00)," 2004, (Last accessed: 08 February 2023).

[26] W. C. Daywitt, "Radiometer Equation and Analysis of Systematic Errors for the NIST Automated Radiometers," National Institute of Standards and Technology, Tech. Rep. 1327, Dec 1983. [Online]. Available: https://doi.org/10.6028/NIST.TN.1327

[27] D. Gu and J. Jargon, "Fundamental Theory of Microwave Measurement; with Applications in Scattering Parameters, Power and Noise," In Preparation.

[28] M. G. Arthur, "The Measurement of Noise Performance Factors: A Metrology Guide," U.S. National Bureau of

Standards, Boulder, CO, Tech. Rep. 142, June 1974.

[29] D. Williams. Microwave Uncertainty Framework. [Online]. Available: https://www.nist.gov/services-resources/software/wafer-calibration-software