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

Todd Schumann, Thao T. Nguyen, Aric W. Sanders
Douglas Boulware, Elyssa Kaplan, Mark Krangle
Joseph R. Mruk, Anthony W. Romaniello, Stephen Segro

NASCTN Chief Engineer: Duncan A. McGillivray
NASCTN Project Manager: Keith A. Hartley
NASCTN Program Manager: Melissa M. Midzor

September 3, 2024

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.

NASCTN provides robust test processes, validated measurements data, and statistical analysis 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.

Representatives from Defense Information Systems Agency (DISA) Program Executive Office (PEO) Spectrum submitted a proposal to NASCTN to measure emissions in the Citizens Broadband Radio Service (CBRS) band to collect data required for DISA PEO Spectrum 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 Capabilities (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 PEO Spectrum 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 Coastal 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 A. Hartley	MITRE	Project management
Mark Krangle	MITRE	Sensor system, and sensor fielding
Elyssa Kaplan	MITRE	Sensor system
Joseph R. Mruk	MITRE	Sensor system

Acknowledgments

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-tiered 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 3-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 the following components:

- 1 - Passive Observation: Characterize aggregate emissions within the CBRS band 3550-3700 MHz in at least two Coastal Dynamic Protection Areas (DPAs), with and without DPA activations. 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 or ground-based DPA (GB-DPA). Assess ecosystem performance to limit CBRS emissions in the GB-DPA.
- 3 - Long Term support - Collection and analysis of Passive Observations throughout sensor deployment, and support transfer of data to a DoD data repository.

This test plan addresses Passive Observation in locations on the east and west coasts in Coastal DPAs.

0.1 Overview of Baseline Measurements and Analysis

In order to expedite sensor deployment and data collection Program Executive Office (PEO) spectrum and NASCTN agreed to work sensor locations, sensor quantities, and deployment of sensors in parallel with Test Plan development. Therefore, there will be noticeable areas where explicit detail cannot be provided in the Test Plan. Any details not available before community review of the Test Plan will be either added to the Test Plan at a later date or documented in the final publicly released test report.

0.1.1 Sensor Deployments

NASCTN has ten (10) sensors built for deployment and data collection in Coastal DPAs between the East and West Coast. The current plan is to deploy five (5) sensors on both the East and West Coasts according to the methodology described in Section 3.3, but the actual distribution of sensors will be decided based on site availability and feasibility. The current installation and deployment efforts for both coasts are as follows:

0.1.1.1 East Coast

NASCTN installed a sensor at Hampton University and another sensor at Norfolk International Terminal. These sites hosted prototypes sensors which have been upgraded to Sensor Rev 1 capabilities.

NASCTN continues to investigate locations outside the Norfolk, VA area to collect aggregate CBRS emissions, which include the Virginal Department of Transportation Chesapeake Bay Bridge Tunnel complex and National Aeronautics and Space Administration (NASA) Wallops Island.

The combination of sensors located near and some distance away from Citizens Broadband Radio Service device (CBSD) deployments will provide a rich and diverse set of CBRS emissions for collection and analysis.

0.1.1.2 West Coast

NASCTN is coordinating with Catalina Island, Midway Museum, Cabrillo National Monument, Miramar, and the San Diego area for potential West Coast sensor installation locations.

The approach is similar to the East Coast to deploy sensors both near and farther away from CBSD deployments to collect both local and aggregate emissions.

0.1.2 Data Collection and Measurements

The sensors will acquire in-phase and quadrature (IQ) data on each 10 MHz channel between 3530 MHz and 3710 MHz in sequence and process it into the data products described in Chapter 5) before discarding the raw IQ data. To increase the frequency of collection, the data processing will occur in parallel with the acquisition from the subsequent channel. The sensors will be configured to take data at a regular interval as quickly as the processing can be reliably completed (100 s). The data will be collected at Institute for Telecommunication Sciences (ITS) servers through the spectrum-characterization and occupancy sensing (SCOS) software where it can be aggregated and analyzed across all sensors.

0.1.3 Data Products

NASCTN sensors will collect IQ emission captures in the 3530-3710 MHz frequency range and reduce the data into the following summary data products to transport to the central repository. The IQ data is discarded after the creation of the data products.

- Power Spectral Density
- Time Domain Power Statistics
- Time Domain Power Summary
- Amplitude Probability Distribution
- Periodic Frame Power

0.2 NASCTN Process Activities

NASCTN projects follow a 5-stage open, transparent, and 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 - NASCTN conducted site surveys for multiple locations in the vicinity of east and west Coastal DPAs and selected several candidate locations.
- Modeling - NASCTN incorporated existing methods and industry provided data to better inform site selections, and to identify measurements to achieve project goals.

- Sharing Ecosystem Assessment (SEA) sensor/ Edge Data collection systems - NASCTN developed data collects and analysis to 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 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 collect emissions in the CBRS band. A number of sensors will be strategically located in the vicinity of Coastal DPAs on the east and west coast to continuously monitor the CBRS environment. 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.

The deliverables planned for this test effort include:

- A test report which includes a repeatable test methodology to collect CBRS ecosystem data to inform on the emission levels within the CBRS band for Coastal 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 to 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

The technical approach and deliverables are anticipated to be refined over time so as to evolve with the changes in the ecosystem in conjunction with needs of stakeholders. Major deviations, though not anticipated at this time, from the outset of the project will be communicated to the sponsor and documented in detail in NASCTN annual technical reports.

Contents

Acronyms

	3GPP	3rd-generation partnership project
	5G	Fifth Generation of cellular mobile communications
250	**AC**	alternating current
	ADC	analog-to-digital converter
	AGL	above ground level
	AMA	Analytical Mechanics Associates
	APD	amplitude probability distribution
255	**ATO**	Authority to Operate
	AWS	Amazon Web Services
	BPF	bandpass filter
	CBRS	Citizens Broadband Radio Service
	CBSD	Citizens Broadband Radio Service device
260	**CCDF**	complementary cumulative distribution function
	CIO	Chief Information Officer
	COTS	commercial off-the-shelf
	CTL	Communications Technology Laboratory
265	**CW**	continuous-wave
	DANL	displayed average noise level
	DFT	discrete Fourier transform
	DISA	Defense Information Systems Agency
	DoD	Department of Defense
270	**DPA**	Dynamic Protection Area
	EIRP	effective isotropic radiated power
	ENR	excess noise ratio
	ENR	excess noise ratio
	ESC	Environmental Sensing Capability
275	**FAD**	full activity dump
	FCC	Federal Communications Commission
	FFT	fast Fourier transform
	FSPL	free space path loss
	GAA	General Authorized Access
280	**GB-DPA**	ground-based DPA
	GEOID	geographic identifier
	GMM	Green Mountain Mesa
	HU	Hampton University
	IJWG	Interagency Joint Working Group
285	**IIR**	infinite impulse-response
	IQ	in-phase and quadrature
	ITM	Irregular Terrain Model
	ITS	Institute for Telecommunication Sciences
	KML	keyhole markup language
290	**LaRC**	Langley Research Center
	LNA	low noise amplifier
	LTE	long-term evolution
	M4	mean, median, maximum, minimum
	MITRE	The MITRE Corporation
295	**Msps**	mega samples per second
	NASA	National Aeronautics and Space Administration
	NASCTN	National Advanced Spectrum and Communications Test Network
300	**NIT**	Norfolk International Terminal
	NEMA	National Electrical Manufacturer's Association
	NF	noise figure
	NIST	National Institute of Standards and Technology
305		
	NSW	network switch
	NTIA	National Telecommunications and Information Administration
	NUC	Next Unit of Computing
310	**P1dB**	1 dB gain compression
	PAL	Priority Access License
	PAPR	Peak-to-Average Power Ratio
	PEO	Program Executive Office
	PFP	periodic frame power
315	**PSD**	power spectral density
	PvT	power versus time
	RF	radio frequency
	RMS	root mean square
	RSA	real-time spectrum analyzer
320	**SAS**	Spectrum Access System

SCOS spectrum-characterization and occupancy sensing

SEA Sharing Ecosystem Assessment

SigMF signal metadata format

SINR signal to interference plus noise ratio

SPDT single pole double throw

SPU signal processing unit

SQL Structured Query Language

SSH secure shell

STIG Security Technical Implementation Guide

SWaP size, weight and power

TDD time division duplex

TEC thermo electric heater/cooler

UE user equipment

UPS uninterruptible power supply

USB universal serial bus

VBA verification bench testing and acceptance

WInnForum Wireless Innovation Forum

List of Figures

List of Tables

Chapter 1

Introduction

Defense Information Systems Agency (DISA) Program Executive Office (PEO) Spectrum proposed a measurement campaign to the National Advanced Spectrum and Communications Test Network (NASCTN) to measure emissions in the Citizens Broadband Radio Service (CBRS) band (i.e., 3550-3700 MHz) to allow Department of Defense (DoD) to collect data required for DISA PEO Spectrum 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 Capabilities (ESCs). The collected data provides DISA PEO Spectrum 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 emissions in the CBRS band in Coastal Dynamic Protection Areas (DPAs) on the east and west coasts through a series of deployed sensors and software to collect and report emissions for visualization and analysis. The test plan describes the methodology, sensor and software design, data collection, modeling, and data visualization/analysis to provide yearly reports and accompanying data on the collected CBRS emissions.

To accelerate the emission collection and generation of data products, PEO Spectrum and NASCTN agreed that the site selection and sensor deployment would occur simultaneously with test plan development. Therefore, there are details of the test plan that are incomplete during the community review and comment period. Any details on the quantity and location of sensor deployment made after completion of this test plan will be documented in updates to the test plan or within the final publicly released test report at the end of the project.

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.

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 3-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: The SAS performs automated spectrum use coordination, and the Environmental Sensing Capability (ESC) detects incumbent use of the band and transmits that information back to the SASs.

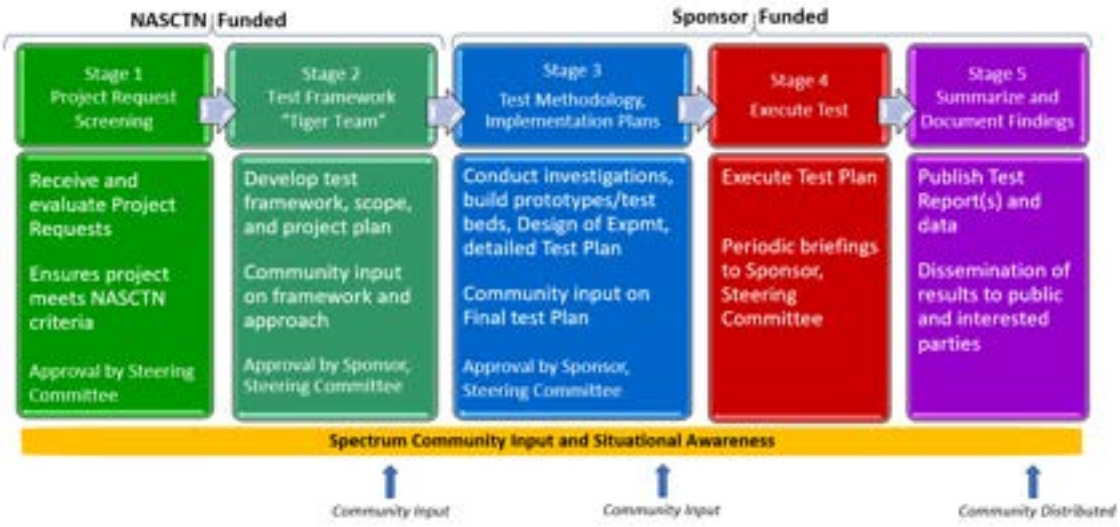

Figure 1.1: NASCTN's 5 Stage Process

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. PEO Spectrum identified one component 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 PEO Spectrum submitted a CBRS Sharing Ecosystem Assessment (SEA) test request to NASCTN to achieve four objectives regarding the CBRS band:

1. Provide data to determine efficacy of permanent sharing between CBRS systems as managed by SAS systems and

ESC systems

- Attempt to obtain corroborating data from CBRS community stakeholders

2. Collect power levels in the CBRS band through continuous automated observations
3. Collect emissions in the CBRS band in the vicinity of San Diego, CA and Norfolk, VA
4. Collect emissions in the CBRS band within at least one Always-On DPA

This DISA PEO Spectrum test request was divided into three major components:

1. Passive Observation: Characterize aggregate emissions within the CBRS band 3550-3700 MHz in the vicinity of at least two Coastal DPAs, with and without DPA activations. 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 or ground-based DPA (GB-DPA). Assess ecosystem performance to limit CBRS emissions in the Always-On DPA.
3. Long Term support - Collection and analysis of Passive Observations over 4 years, and support transfer of data to a DoD data repository.

The project described in this test plan addresses item 1 (Passive Observation: Characterize aggregate emissions within the CBRS band with and without DPA activations in the vicinity of at least two Coastal DPAs) along with the collection and analysis of the data over approximately 18 months. A separate test plan was developed and available on the NASCTN website for item 2 and the associate data collection and analysis.

Chapter 2

Data Collection Strategy Overview

The initial test request sent to National Advanced Spectrum and Communications Test Network (NASCTN) identified data capable of validating and verifying adherence to the Federal Communications Commission (FCC) Part 96 rules as they pertain to Citizens Broadband Radio Service (CBRS) operations in the shared band. This request was originally examined and it was determined that certain of the Part 96 rules could be validated, but only with additional and corroboration data from other systems, such as activation logs from Environmental Sensing Capabilities (ESCs). Through the project, it was not possible to obtain such corroborating data, so the goals of the project instead focus on taking data which is of interest to the entire stakeholder community and can be used to assess the effectiveness of the ecosystem, though not directly validate its adherence to the Part 96 rules.

For example, to completely assess the viability of the CBRS Sharing Ecosystem Assessment (SEA) ecosystem, a measurement system able to measure to -144 dBm/10 MHz is required as well as deploying sensors at multiple points inside a Dynamic Protection Area (DPA). Neither of these criteria is possible or realistic. Instead, NASCTN developed a sensor that measures to approximately -100 dBm/10 MHz to deploy in various locations outside of a Coastal DPA to collect useful emission data and provide a variety of data products to the CBRS stakeholder community that can give insights to the behavior of the ecosystem as a whole.

NASCTN will collect real-world radio frequency (RF) data products in the CBRS band to help inform stakeholders on the efficacy of the ecosystem as it's deployed. The high-level collection strategy described here will serve as a template to what the community can expect from the project based on the design decisions, data captured, and final outputs and reports. Further detail on each of the sections below can be found in later chapters of the text.

2.1 Site Selection

As of this writing, NASCTN has built ten (10) sensors for deployment in the vicinity of multiple DPAs on the East Coast and West Coast. The plan is to deploy five (5) sensors on the East Coast, in the Norfolk, VA area and five (5) sensors on the West Coast in the San Diego, CA area. The sensor locations are balanced between ideal and practical locations available for installation. Modeling combined with physical site surveys were used to provide confidence that the selected locations will be able to observe emissions from both incumbent and CBRS deployments. This will allow data products that observe the DPA with and without activations for at least two DPAs, one on the East Coast and one on the West Coast.

Further, sites were chosen based on the anticipated usefulness from the data collected from them. Several current deployments as well as anticipated deployments are mentioned in the text. Sensors will be split between measuring from within a cluster, or dense deployment, of Citizens Broadband Radio Service devices (CBSDs) in a metropolitan area and a sensor outside the metropolitan area observing back in. The former will allow measurement of the local dynamics of the CBRS ecosystem as well as validation of actions of the system. However, these will likely be limited to measuring the nearest-neighbor CBSD due to the large difference in propagation loss over relatively short distances. For this reason, a sensor will also be located outside the CBSD cluster, looking back in with no local CBSDs around

it (to the extent possible). These sensors will measure the aggregate effects and behavior of the system, including aggregation of CBSD energy. Finally, NASCTN will place sensors in other locations of interest, such as Camp Pendleton, which will provide useful data about the ecosystem and lend insights to the overall behavior and operation.

Chapter 3 provides a more detailed overview of the site selection process, including the iteration between modelling of selected sites and physical site surveys to narrow the selection of feasible locations for each of the sensor objectives described. In addition, it provides a summary of the current sites under consideration and the reasoning/usefulness of each site.

2.2 Sensor Design

The sensor was originally optimized to reduce the noise floor as much as possible. Early measurements of the local CBRS deployments indicated that actual observation of emissions in the CBRS band would be challenging and require very sensitive sensors. Additionally, a sensor with as low of a noise floor as possible would be needed to observe emissions within DPAs when activated and nearby CBSDs are vacated.

In practice, this has shown to be limiting as many of the ambient signals in the CBRS ecosystem are higher in power, especially as deployment densified. The sensors have been modified to account for these higher-level emissions through both protection of adjacent channel overloads as well as through re-optimization of the sensor to a better balance between dynamic range and sensitivity. The sensor can be dynamically adjusted in the field without manual intervention in the case that its surrounding environment changes (e.g., new deployments of CBSDs require re-tuning of the dynamic range settings).

When deployed, the sensors are first characterized in a variety of dynamic range settings to characterize the overload levels from surrounding emissions. The sensor is then optimized to have the lowest noise floor while still protecting from the majority of overloaded measurements from the nearby CBSDs, which allows measurements both during DPA activations as well as when the DPA is not activated. The reduction in sensitivity from these measurements was deemed prudent to collect meaningful data throughout the project on the entire ecosystem, both with and without DPA activations.

The incoming data will be analyzed on a periodic basis to ensure that the sensors are still within this ideal dynamic range. If the deployment landscape around the sensor changes, the dynamic range settings of the sensor will be adjusted to keep the sensor with the lowest noise floor, but not subject to a high percentage of overloaded measurements.

The bulk of the sensor remains the same as in the previous NASCTN CBRS SEA Test Plan for Always-on DPAs [2] with minor improvements to the RF input to help mitigate overloads. The backend was altered to improve the size, weight and power (SWaP) of the sensor and allow a ruggedized, weather-hardened design. This allows more flexibility in deployment locations as the sensor is no longer tied to a split-system approach. The sensor will still communicate with the central servers through cellular backhaul to also allow flexibility in deployment for locations where network connectivity is not guaranteed.

Chapter 4 provides a detailed description of the sensor build and software as it differs from the previous Test Plan [2] as well as the reasoning for each change. In addition, to allow full documentation of the sensor design in the case that additional sensors may be built, Appendix A describes the complete sensor build including all detail originally present in the prior Test Plan [2].

2.3 Sensor Data

The data products developed for collection are reduced from raw in-phase and quadrature (IQ) captures. Raw IQ data presents a security and privacy concern and is otherwise impractical to record on a near-continuous basis due to

its extremely large data size. The reduced data products (described in more detail in Chapter 5) provide amplitude statistics and spectral behavior that is useful to the community at large. They incorporate both measurements across the entire 10 MHz channels, both in frequency and time, as well as with finer spectral resolution. These statistics will prove useful in determining the efficacy of the overall ecosystem as well as giving insights to the system behavior such as aggregation of signals, short term-dynamics, long-term trends, etc. Preliminary examples in using these data products to characterize the behavior of the ecosystem can be found in Chapter 6.

The data products are collected from the sensors as they are generated and imported to an Structured Query Language (SQL) database. Given the fixed tabling of the database, changes to the data products become extremely arduous, both in incorporating changes or alterations to the data products as well as the addition of new data product elements. For this reason, the data from the prototype deployments collected over the previous year was analyzed to determine gaps and additions that would be useful before fixing the data products. Alterations to the data products are technically possible after finalizing the database schema, but would have to show extreme benefit to be incorporated.

2.4 Deliverables

The collection of data products will be delivered to the community publicly at the end of the project (currently projected to be the Fall of 2025) accompanied with a report summarizing the data and high-level findings. The report will also detail any non-idealities in the data, any deviations from this test plan, and any considerations that need to be taken into account when using the data, such as interpretation of overloaded measurements. There will be limited NASCTN-run analyses in the report, instead focusing on example analyses to serve as templates for commercial and government stakeholders to aid in both understanding and using the collected data. Several of these analyses are detailed in Chapter 6 below. The stakeholder community can use the collected data to run their own independent analyses to showcase the behavior of the ecosystem, which will lead to improvements to the system over time and improve the efficacy of the CBRS sharing and model.

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 San Diego, CA and Norfolk, VA. Given the vast geographical areas of interest, modeling is used as a tool to help inform site selections for NASCTN SEA sensors to measure emissions in the CBRS band.

The NASCTN modeling leverages the industry provided obfuscated full activity dump (FAD) data on Citizens Broadband Radio Service device (CBSD) deployment and the Wireless Innovation Forum (WInnForum) reference Dynamic Protection Area (DPA) move list algorithm to estimate the aggregate emissions at potential sensor locations when the DPA is either active or inactive. The modeling result provides an estimate, typically an upper bound, of the aggregate power level received within any 10 MHz channel in the CBRS band at a specific sensor location. To assess the accuracy of the modeling results, the estimated aggregate emissions are compared with measurements collected at several NASCTN prototype sensor locations. Quantifying the offset between the modeling results and measurements provides insights into key technical challenges and assumptions with the current modeling, and sets boundary for aggregate emission estimates at other potential sensor locations.

Once areas of interest are modeled, predicting anticipated CBSD emission levels, physical site surveys are conducted in the areas at specific locations which indicate that a placed sensor would be able to achieve its desired function (measure local-dynamics, observe aggregate emissions, etc.). As the placement of the sensors is ultimately tied to feasibility of deployment, deployable locations will be a limited subset of the modeled locations. The site modeling and physical survey process is iterated to find the best suitable locations for the sensors in a given area.

3.1 Overview of Modeling and Simulation

3.1.1 Obfuscated FAD Data on CBSD 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 on a quarterly basis [3]. In a joint collaboration to support the SEA project, NASCTN received a subset of the obfuscated FAD data for selected areas from NTIA ITS. These include two areas near DPA West-14 and DPA East-1 used for Task 1 test plan development. 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 [4].

Details of how the SAS administrators build the obfuscated data from FAD data can be found in [3]. In this project, the obfuscated data were utilized to generate randomized CBSD deployments for each area. Three categories of CBSDs, i.e., indoor Category A CBSDs, outdoor Category A CBSDs, and outdoor Category B CBSDs, were generated separately for each census block geographic identifier (GEOID) provided in the obfuscated data file. And for each category, the CBSD registration request data and grant request data were randomly generated from the obfuscated

distributions, and subsequently the grants were randomly assigned to the CBSDs. Note that in the CBRS standards [5], 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.

3.1.2 WInnForum Reference DPA Move List Algorithm

Regulatory rules [6] 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 specified 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 [7, R2-SGN-24] and [8].

The primary output of the move list algorithm is the move list for each 10 MHz channel for each DPA. 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. Figure 3.1 shows example distributions of move list size and keep list size for ten channels for two representative West-14 and East-1 DPAs using obfuscated FAD data generated on 01 January 2024.

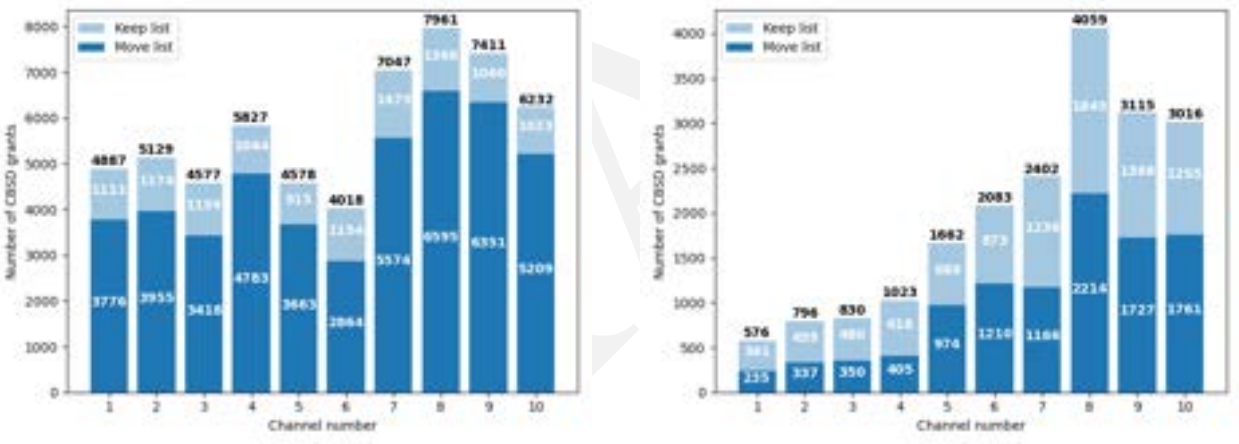

(a) DPA East-1

(b) DPA West-14

Figure 3.1: Distributions of move list and keep list sizes for areas of interest.

As of this writing, the CBRS Interagency Joint Working Group (IJWG) has proposed several changes that significantly impact the DPA protection calculations and results. The changes include:

- *Duty cycle and network loading factor*: Apply 8 dB reduction for duty cycle and network loading factor to all CBSDs when calculating path loss or interference contribution to any DPA protection point,
- *Clutter*: Apply median clutter loss (from ITU-R P.2108, section 3.2) [9] to all CBSDs that have an antenna height $\leq$ 6 m above ground level (AGL) when calculating path loss or interference contribution to any DPA protection point,
- *DPA neighborhoods*: Use 6 separate neighborhood distances for Category B CBSD ($>$ 6 m AGL), Category B CBSD ($\leq$ 6 m AGL), outdoor Category A CBSD ($>$ 6 m AGL), outdoor Category A CBSD ($\leq$ 6 m AGL), indoor Category A CBSD ($>$ 6 m AGL), and indoor Category A CBSD ($\leq$ 6 m AGL),

- *Path loss and interference calculation*: Use a fixed 50 % reliability (median) for Irregular Terrain Model (ITM) path loss and interference calculation, rather than a 95th percentile aggregate emission from a 2000-iteration Monte Carlo with randomized ITM reliability values ranging from 1 % to 99.9 %.

To assess the impacts to move list calculation, NASCTN has implemented these DPA changes and re-run the simulation with the same DPA configuration and CBSD deployment. Figure 3.2 shows distributions of move list size and keep list size for West-14 and East-1 DPAs taking into account the DPA changes.

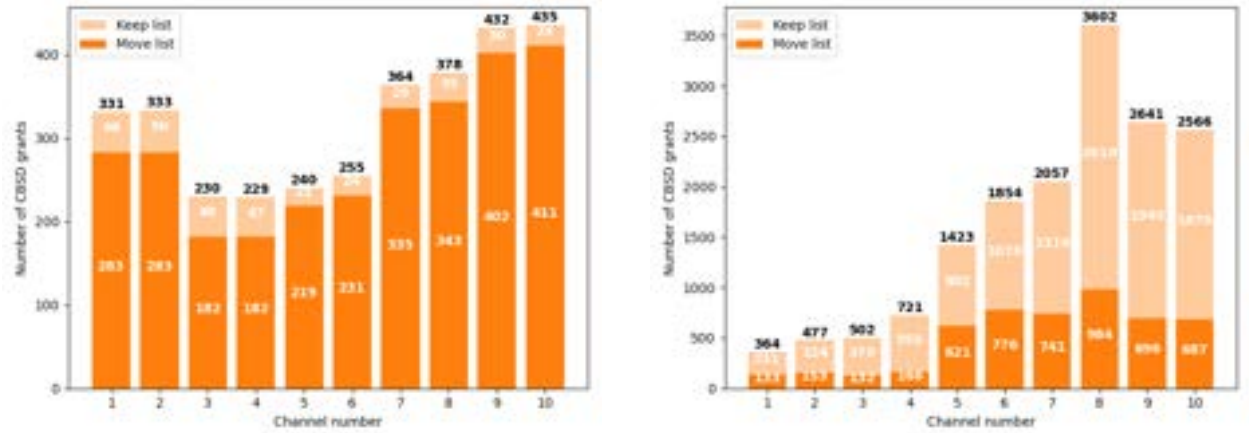

(a) DPA East-1

(b) DPA West-14

Figure 3.2: Distributions of move list and keep list sizes for areas of interest with implementation of proposed DPA changes.

With the proposed DPA changes, the neighbor list size is decreased due to smaller DPA neighborhood distances. In addition, the move list size is also decreased due to other factors including 8 dB effective isotropic radiated power (EIRP) reduction, clutter, and ITM median loss.

3.1.3 Aggregate Emission Calculation

This section describes the aggregate emission calculation at potential sensor locations in the vicinity of DPAs. The calculation is a modified version of the WInnForum aggregate emission check at protection points inside DPAs [7, R2-SGN-12] and [8]. Instead of using a reference 1D antenna described in [7, R2-IPM-04] at the protection point, more accurate representative 2D antenna patterns are used in the modeling.

Given a SEA sensor location s and any 10 MHz channel ch in the 3550-3650 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 neighborhood around a sensor is defined as the area in which transmission of a CBSD grant can be detected by the sensor. In the current simulation, the sensor's neighborhood is considered the same as the monitoring DPA's neighborhood. The model then computes the emission contribution $I_{i\ (\mathrm{dBm})}$ from each individual grant i, for $1 \leq i \leq N$, to the sensor s. Assuming coherent transmission among the grants, 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 $I_{i\ (\mathrm{dBm})}$ of an individual grant i, the model uses a simple link budget equation as follows:

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

where $EIRP_i$ is the 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).

To calculate the aggregate emission 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). Note that the aggregate calculation assumes coherent transmissions of CBSDs within the channel ch.

Figure 3.3 depicts protection points (yellow markers) inside DPAs East-1 and West-1 and potential sensor locations (green markers) nearby these DPAs. The aggregate emission at the protection points and potential sensor sites can be estimated using the calculation described above.

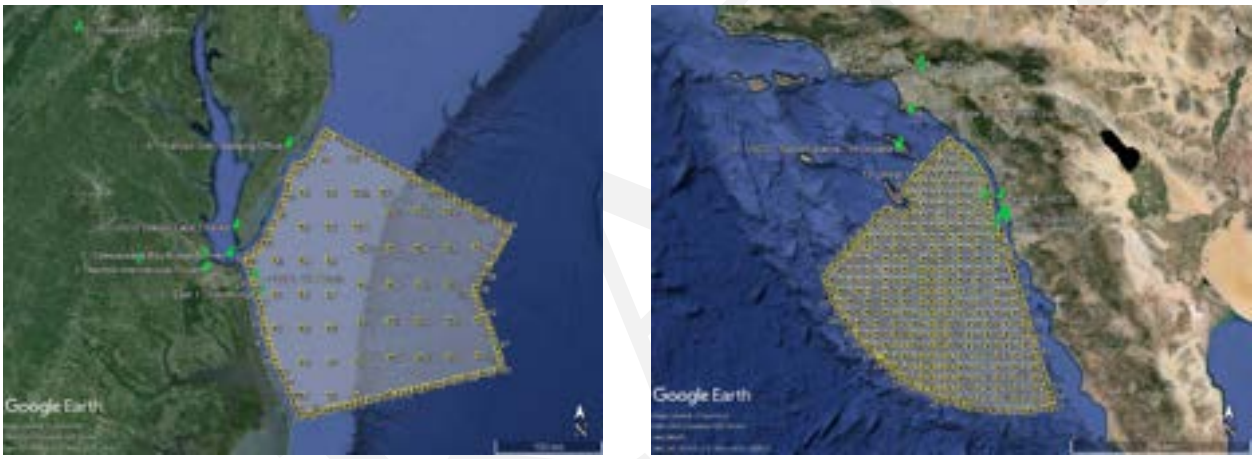

(a) DPA East-1 (b) DPA West-14

Figure 3.3: Illustrations of protection points and potential sensor locations in the vicinity of DPAs East-1 and West-14.

Without any DPA changes, Figure 3.4 shows examples of aggregate emission estimates on (3590–3600) MHz channel at different locations as follows:

- Protection points inside DPA East-1 (Figure 3.4a),
- Protection points inside DPA West-14 (Figure 3.4b),
- Potential sensor locations near DPA East-1 (Figure 3.4c),
- Potential sensor locations near DPA West-14 (Figure 3.4d).

At each location, the aggregate emission is estimated for both DPA inactive and active states. Note that 1D reference antenna pattern [7] is used for calculation at protection points, whereas 2D omnidirectional antenna is used for estimation at sensor locations.

With the implementation of DPA changes, the impacts to the aggregate emission estimates at protection points and potential sensor locations in the vicinity of DPAs East-1 and West-14 can be examined in Figure 3.5. When the DPA is inactive, the results at both protection points and sensor locations in Figure 3.5 have decreased by 8 dB or more as

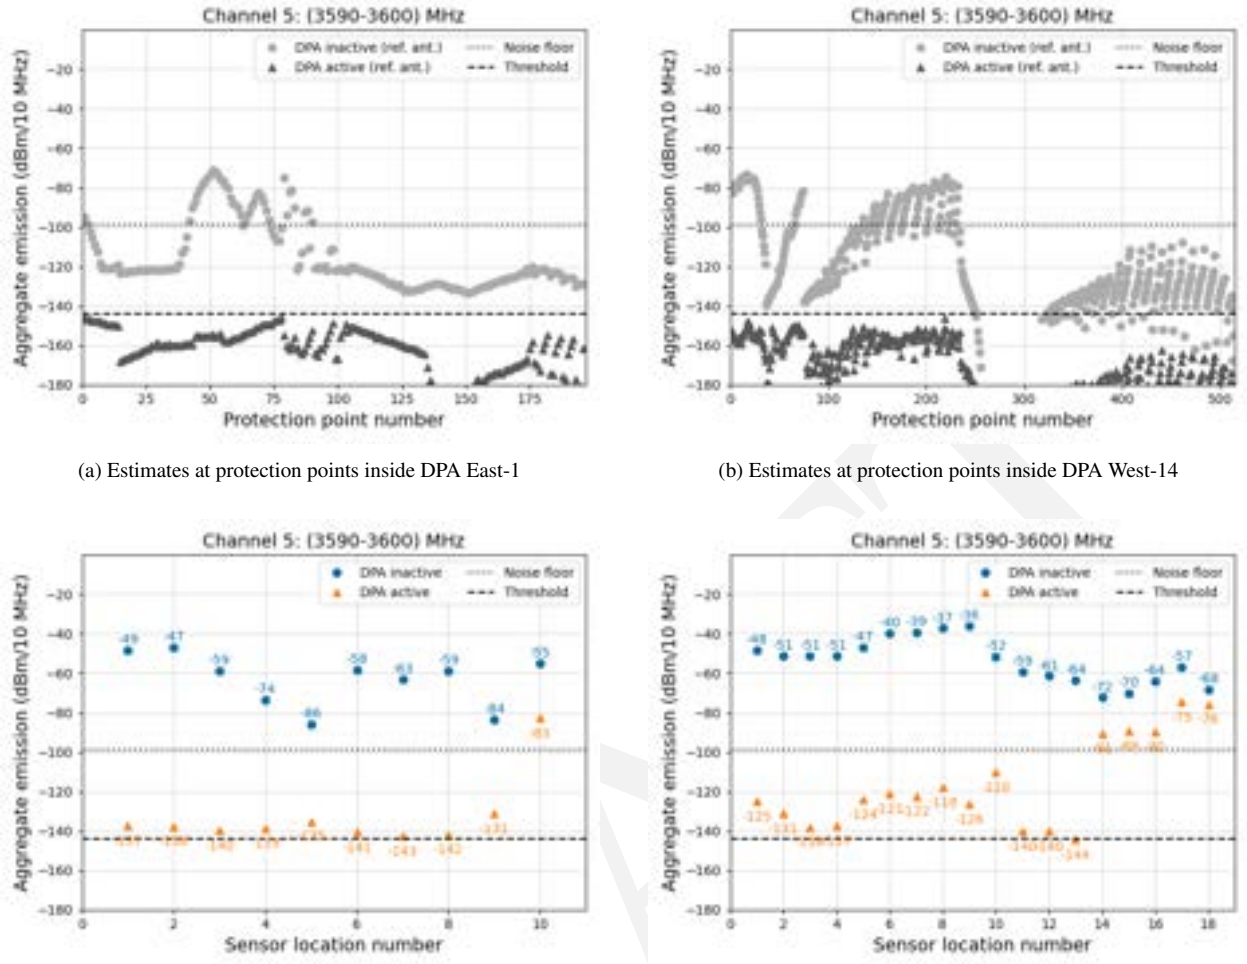

(a) Estimates at protection points inside DPA East-1

(b) Estimates at protection points inside DPA West-14

(c) Estimates at sensor locations near DPA East-1

(d) Estimates at sensor locations near DPA West-14

Figure 3.4: Examples of aggregate emission estimates on (3590-3600) MHz channel at protection points and potential sensor locations in the vicinity of DPAs East-1 and West-14 when the DPA is active and inactive prior DPA changes.

compared to the results in Figure 3.4. However, when the DPA is active, the aggregate emissions at protection points stay below the protection threshold (black dashed line), but the emissions at several sensor locations inside the new DPA neighborhood have increased due to more CBSDs closer to these locations are allowed to stay on the keep list. For those sensor locations outside the new DPA neighborhood (e.g., sensor locations 9 and 10 in the vicinity of DPA East-1), the emissions have reduced due to the lack of CBSDs outside the neighborhood.

3.2 Comparison to Prototype Sensor Data

To evaluate the modeling results for both prior and post DPA changes, real measurements collected at Hampton University (HU) and Norfolk International Terminal (NIT) prototype sensor locations on 01 January 2024 are utilized. Figure 3.6 shows the comparison between modeling results and prototype data at HU location (upper row) and NIT location (lower row). The figures in the first column and second column are modeling results prior and post DPA changes, respectively. Each of these figures shows the aggregate emission estimated for every 10 MHz channel in (3550–3650) MHz band. The blue dots represent the aggregate emissions when the DPA is inactive, whereas, the

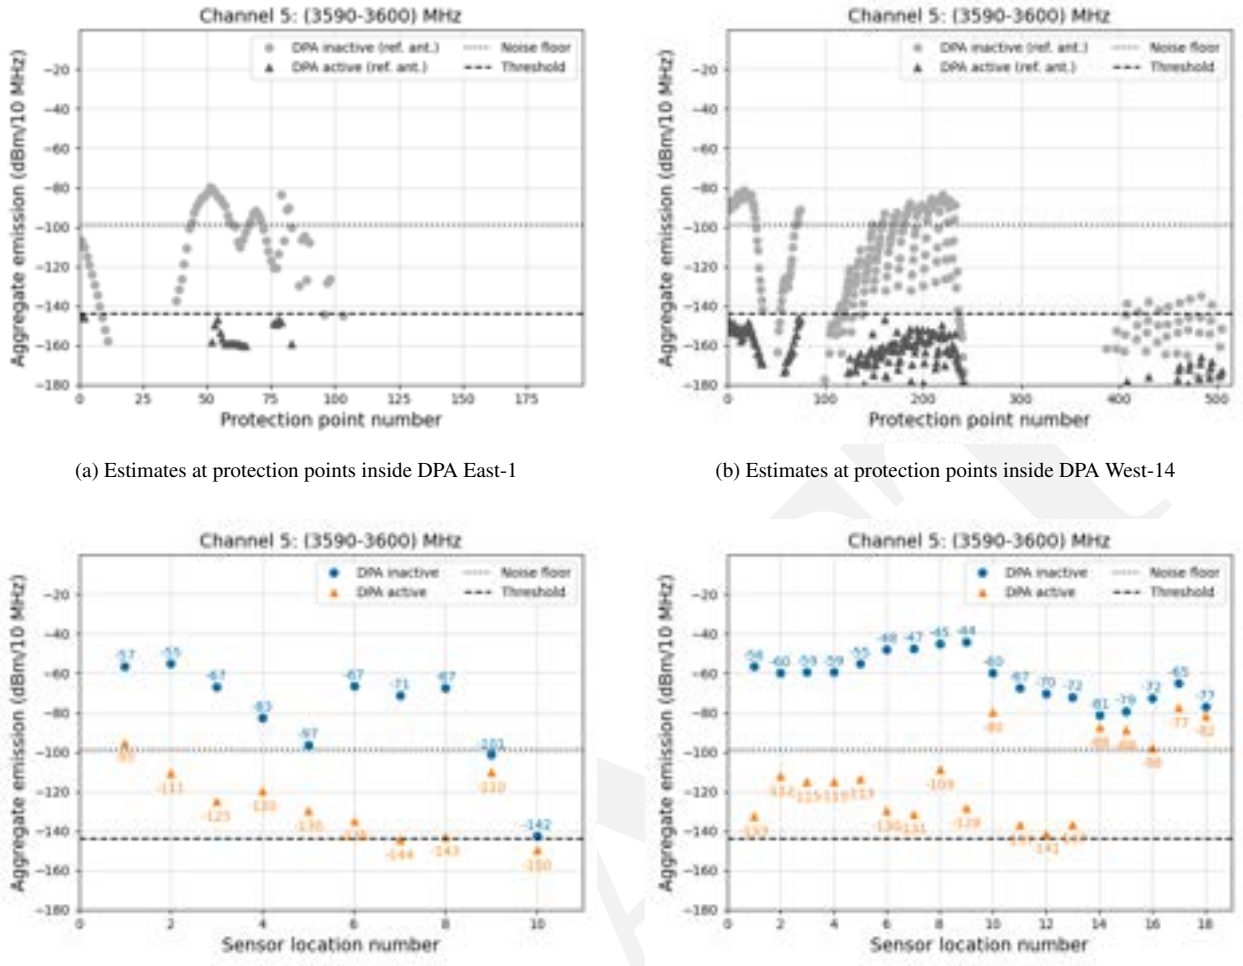

(a) Estimates at protection points inside DPA East-1

(b) Estimates at protection points inside DPA West-14

(c) Estimates at sensor locations near DPA East-1

(d) Estimates at sensor locations near DPA West-14

Figure 3.5: Examples of aggregate emission estimates on (3590-3600) MHz channel at protection points and potential sensor locations in the vicinity of DPAs East-1 and West-14 when the DPA is active and inactive with implementation of proposed DPA changes.

orange triangles represent the emissions when the DPA is active. The aggregate emission variations are due to a 10-iteration Monte Carlo of CBSD deployments randomized from the 01 January 2024 obfuscated FAD data. The figures in the last column show the measurements collected in each channel during that day. The yellow dots and green dots represent max peak periodic frame power (PFP) measurements and mean root mean square (RMS) PFP measurements, respectively, collected at different times during the day.

Except for some of the upper channels in which the prototype sensors may experience overload conditions, the lower channels provide more reliable measurements at both HU and NIT locations. Since there was no DPA activation observed on this day, only modeling results for DPA inactive state are compared with max peak PFP measurements. As shown in Figure 3.6, the measurements in the lower channels are smaller than the results of prior DPA changes, and they are equal to or slightly larger than the results of post DPA changes. Examination of the proposed DPA changes and their impacts to the modeling results are still on going.

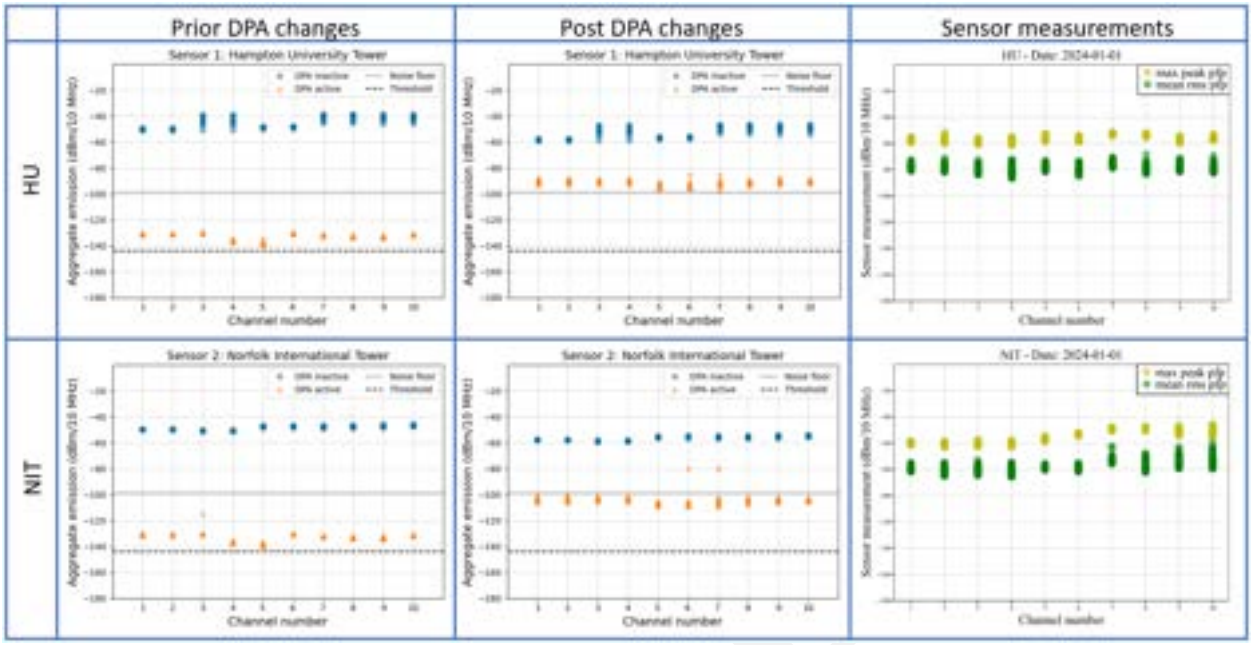

Figure 3.6: Modeling results with prior and post DPA changes vs. measurements collected at HU and NIT prototype sensor locations.

3.3 SEA Sensor Siting

In discussions with MITRE representatives to the Department of Defense (DoD) Chief Information Officer (CIO), an overall deployment strategy was devised for sensors in a given area. In metropolitan areas, there are dense CBSD deployments as indicated in the obfuscated FAD data. Because of the relatively short distances to neighboring CBSDs, the path loss to the sensor can differ drastically, even for CBSDs that differ only slightly in relative distance radially from the sensor. The result of this is that a sensor deployed within a CBSD cluster in a metropolitan area tends to dominantly measure only the nearest-neighbor CBSD. This can be seen in the data taken from the prototype sensors, where it appears as though only a single CBSD is measured.

Investigating path loss curves (taking free space path loss (FSPL) as an example), these curves typically "flatten" with increasing radial distance due to the overarching logarithmic trend. The result is that CBSDs that have similar distances from the sensor experience a smaller delta in path loss the further away from the sensor they are located. Within the CBSD cluster where the separations are relatively small, this leads to a dominant CBSD in most cases. However, if the sensor is located further away from the CBSD cluster, it would be more able to measure true aggregate emissions from the cluster.

Therefore, it is proposed to locate sensors both within and outside CBSD clusters in general metropolitan areas (or other areas of denser CBSD deployments). Taking the prototype sensor deployments in the Norfolk, VA region, the two deployed sensors at NIT and HU are both within the CBSD cluster. These serve to measure the local CBSD response and provide validation of the events within a DPA neighborhood. In addition to these sensors, an additional sensor located outside the metropolitan area would allow for an aggregate view of the same CBSD cluster. In the case of the Norfolk, VA region, Cape Charles or a location nearby is identified as an ideal spot to deploy an aggregate viewing sensor. This location has a vantage point towards the Norfolk region, but without deployed CBSDs nearby, which would allow it to measure aggregate emissions in the Norfolk area without being dominated by a nearest-neighbor CBSD.

To aid in determining ideal locations, a sensor radio frequency (RF) horizon mapping tool was developed, which calculates the sensors' viewing horizon based on the measured sensor properties (e.g., noise figure), deployment parameters (e.g. antenna mounting height), and the ITM propagation model. As sites are identified, this tool, with an overlaid example CBSD laydown generated from the obfuscated FAD data, gives insight into what the sensors expect to observe from a given deployment location. As shown in Figure 3.7a, an example of a generated sensor horizon from the Cape Charles area, it can be seen that the Norfolk area is at the outer edge of the sensor horizon. In this case, it may be prudent to use a more directional antenna to expand the sensor horizon in the direction of interest. If a directional antenna is used, it should be noted that the overall aggregation of energy into the antenna will differ from an omni-directional antenna used elsewhere, which should be considered in analysis (note that the individual data collections identify the antenna used by the sensor as well as the pointed direction).

Based on these tools, site surveys have been conducted with an interest both in observing CBSD emissions, but also in terms of their ability to observe aggregate emissions from multiple CBSDs distanced away from the sensor as shown in Figure 3.7b. As locations are narrowed through the modeling and physical surveys, final sensor placements are determined based both on feasibility of deployment as well as how well measurements taken at the site of interest match the specific objectives of the sensor (e.g. measuring aggregate emissions without being dominated by a nearby CBSD).

With the number of sensors currently produced (ten), it is feasible to deploy two sets of sensors measuring in the above configuration, one near the Norfolk, VA region and one in the San Diego, CA region. The remaining sensors will be sited based on additional goals, such as ground-based DPA (GB-DPA) behavior at Camp Pendleton or additional aggregate views from Catalina Island. Several analyses are being performed to indicate the needed number of sensors to be deployed in the vicinity of a DPA for each specific measurement objective (e.g. aggregate views, model verification, etc.), noting that the required number will likely exceed NASCTN's ability to produce sensors. Details of these analyses will be included in the final report to allow other stakeholders to continue the effort with additional sensors if desired.

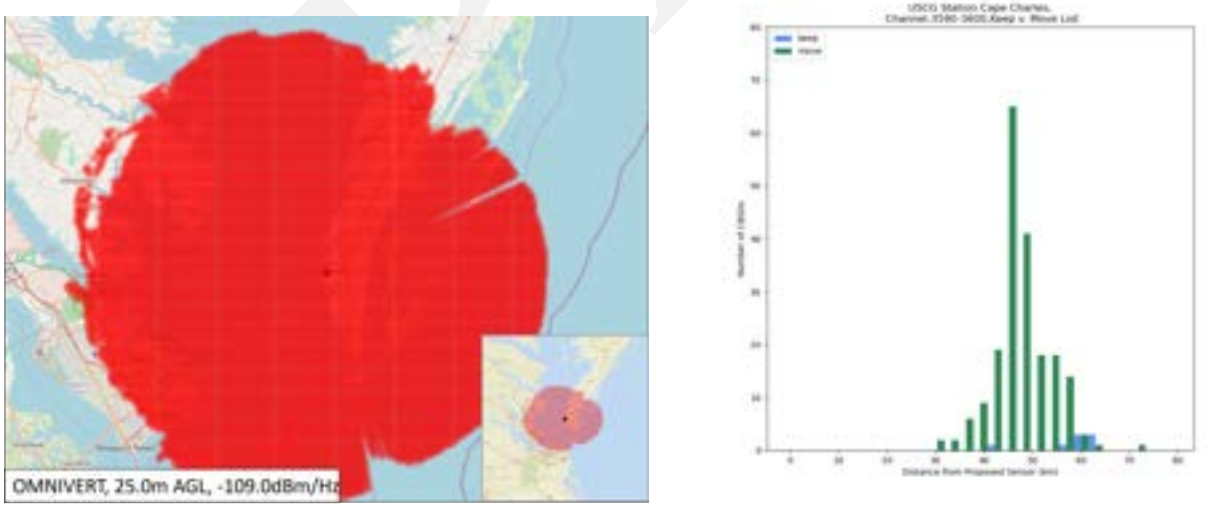

(a) Sensor RF horizon mapping

(b) Histograms of move and keep lists on (3550–3600) MHz

Figure 3.7: Example of sensor RF horizon mapping and CBSD move and keep list histograms at Cape Charles location.

3.4 Site Survey

In addition to using modeling as a tool to inform site selections, NASCTN has been actively conducting site surveys for potential sensor locations in the vicinity of DPAs East-1 and West-14.

3.4.1 Camp Pendleton and West-14

NASCTN revisited Camp Pendleton and the surrounding areas near the West 14 DPA to investigate possible deployment locations on the West coast. While Camp Pendleton is itself a GB-DPA, it is also positioned to observe the West-14 DPA. As the deployment location within Camp Pendleton is nearing certainty, a final survey was performed at the location, including utilizing an antenna with a higher mount point as would be deployed to confirm the expected emissions in the area. Based on this survey, a final installation scheme was developed to deploy the sensor.

In addition to Camp Pendleton, NASCTN also surveyed locations ranging along the coast line of California from Long Beach to San Diego. It was observed that there were numerous CBRS deployments all along the coast at every location surveyed. Of particular interest was San Diego, which serves as a CBSD cluster similar to the Norfolk, VA region. A general strategy was devised similar to East-1 where sensors could be deployed within the San Diego area measuring the local emissions and an additional sensor could be deployed outside observing the aggregate emissions. Due to the terrain to the east of San Diego, no suitable locations for a sensor could be found further inland. However, along Point Loma is promising with few CBSDs deployed nearby and relatively similar distance profile to the San Diego CBSDs where the path loss curve begins to "flatten." A subsequent site survey to points along Point Loma is planned to observe the expected emissions from that vantage point.

3.4.2 East Coast

Further site surveys have been conducted and are planned in the vicinity of the East 1 DPA. As seen above, the Cape Charles area is well positioned to observe an aggregate view of the Norfolk area where the prototype sensors are deployed. A previous site survey to the area indicated that emissions from the Norfolk area were not visible from Cape Charles using the same omni-directional antenna as the current sensors are deployed with, which matches with the sensor horizon mapping tool. A further study of this location is planned where a higher gain, directional antenna will be used to increase visibility in the direction of Norfolk, VA. In addition, locations closer to the city are being investigated, such as the Chesapeake Bay Bridge Tunnel entrances as possible locations in the case that the directional antenna does not give the needed range on the coast of Cape Charles.

The neighborhood of the East 1 DPA stretches far inland beyond the coastal Norfolk, VA region. Because of this, sensor siting is also taking place in more inland regions. Two regions of interest based on the obfuscated FAD data are Washington, D. C. and Richmond, VA. An initial survey was conducted in the Washington, D. C. area at the MITRE facility in McLean. Results from here showed that there is a strong CBRS presence indicating it would be a good location for a sensor deployed within the CBSD cluster. In both of these locations, however, it has proven difficult to find a suitable location for deployment that would give an aggregate view of the CBSD cluster without a nearby dominating CBSD. Additional locations in both these areas are being further investigated to ascertain if an aggregate view of either can be obtained. In addition, the modelling is being used to determine if these inland sites would still be of interest if the proposed changes to the move-list algorithm are adopted.

3.4.3 Catalina Island

A site survey was also performed in the channel island of Catalina Island. Based on the obfuscated FAD data, there are few CBSDs deployed on the island and a sensor installed on the island would give a relatively unobstructed, over-water view of the mainland, with a relatively high radial distance to the nearest CBSDs allowing a more aggregate view.

When surveying the location at higher elevations, the survey system was modified to collect the same data products as the deployed sensors. This coupled with the similar noise figure between the two systems allows the surveys to collect very similar data as a sensor deployed in the same area. Of particular interest in the Catalina Island survey were the PFP plots. These give evidence of measured aggregate emissions from CBSDs on the mainland, indicated by the distinguishable power levels as well as the uplink portions of the time division duplex (TDD) split being filled in with additional emissions. Further analysis will be conducted to determine whether and how much these signals are aggregating.

In addition, the survey was used to showcase emissions taken with a directional and omni-directional antenna. The California coast line parallels the island, giving a view at a relatively wide angle. The directional antenna with a beamwidth of 15° was used to collect emissions in radials from the island towards the mainland. In addition, an omni-directional antenna was used to collect emissions with a lower gain from the entirety of the coastline. The two measurements show different levels of apparent aggregation, with presumably fewer and higher power CBSDs with the directional antenna. This is also evident with the omni-directional antenna collecting emissions with a higher power than the directional antenna despite the lower peak antenna gain.

Based on these observations, a sensor deployed on Catalina Island at a high elevation to observe aggregation of CBSDs is of high interest. While technically not within the West-13 DPA, it is entirely surrounded by the DPA. Given the method of calculating emissions to protect a DPA, when the West-13 DPA is activated, it is likely that a sensor deployed on Catalina Island would be protected to the same threshold as a protection point at the same location, offering a view essentially from within the DPA. In addition, the path to the mainland DPAs is predominantly an over-water path collecting propagation effects unique to this type of propagation. Finally, with few nearby CBSDs, a sensor deployed here would have a good chance of observing aggregate emissions from the mainland. In this vein, it is also being investigated whether a differential study between the aggregation from both an omni-directional and directional antenna would be of interest to study how CBSDs tend to aggregate.

Chapter 4

Sensor Architecture

The overall sensor architecture remains similar to that of the previous Test Plan [2, Chapter 4] with some minor modifications. The complete sensor architecture is documented in Appendix A to give a complete view of its construction, while the modifications and reasonings for these changes are summarized in this chapter.

4.1 Sensor Design

Site selection has been a limiting factor in the sensor deployments. To increase the viability of sites, the sensor signal processing unit (SPU) has been redesigned to be ruggedized and weatherproof, leading to an all-outdoor sensor that can be deployed with only the need of 120 VAC power and commercial cellular coverage. The SPU is fitted with a thermo electric heater/cooler (TEC) to maintain the internal components within their specified operational temperature ranges. Additionally, the size and weight of the SPU were greatly reduced to allow for more flexible maneuvering and mounting options (such as carrying it through an attic door or mounting it on a roof). The SPU is equipped with a custom length power cord to reach to a nearby power source. With these changes, there are now options for a split deployment (preselector and SPU are in different locations and connected through a long cable harness) or a combined deployment (preselector and SPU are collocated with a short cable harness) for the sensor. The combined deployment eliminates the requirement of having to run RF and network cables from the preselector through buildings to an internally located SPU and increases the viability of sites that cannot support the split deployment. The new sensor design is currently being field tested in Boulder, CO to ensure it can operate fully outdoors.

The preselector band-pass filter was redesigned to shift the upper and lower cutoff frequencies lower in frequency. Observations from the prototype sensors showed that high power emissions from the adjacent channel C-band prevented the ability to measure the upper 2-3 General Authorized Access (GAA) CBRS channels (3670-3700 MHz). The signal analyzer collects data within a fixed 56 MHz bandwidth and then filters and down-samples to the desired 10 MHz acquisition. Because of this, the analog-to-digital converter (ADC) is open to a 56 MHz window encompassing the desired measurement channel, but is subject to overloads from signals within the entirety of the window. In the upper CBRS channels, this window extends into the C-band above CBRS, which has a much higher EIRP limit than CBRS. Because of this, the deployed sensors frequently experienced overload conditions when collecting data in these upper CBRS channels, thus preventing meaningful data from being collected. To help counteract this, the band-pass filter was redesigned to have a roll-off starting before 3700 MHz to significantly improve the attenuation of the adjacent C-band at the cost of reduced sensitivity in the upper CBRS channels. This will eliminate most of the overload effects to the signal analyzer's ADC caused by C-band emissions and allow the upper channels of CBRS to be calibrated and measured with confidence that ADC clipping has not impacted the results. The revised filter design also extends the passband further into the 3.45 GHz band, allowing a preliminary observation into emissions below 3550 MHz. With these revisions, the measurement band is extended to cover the span: 3530–3710 MHz.

In a similar vein, data collected from the prototype sensors have shown that CBRS channels may become overloaded within the measured band if a CBSD is sufficiently close or if a new CBSD deployment occurs nearby. These in-band overloads cannot be mitigated with the new front-end bandpass filter. In these cases, the dynamic range settings of the

signal analyzer, namely the signal analyzer attenuation and reference level, are adjusted on a per-channel basis such that the channel can still be measured. This degrades the noise figure within a given channel, but allows the channel to be calibrated and measured without the effects of ADC clipping. Since the signal analyzer samples on a 56 MHz window, it is also required to adjust the dynamic range settings for adjacent channels as well.

Additionally, a sampling approach is being investigated which could leverage the hardware anti-aliasing filter in the signal analyzer that exists before the ADC. In this approach, a full 40 MHz capture is performed with the desired channel being either the upper or lower 10 MHz. In this case, the anti-aliasing filter provides additional protection from adjacent-channel emissions, either above or below the measured channel. The recorded in-phase and quadrature (IQ) waveform is digitally frequency shifted such that the desired channel is centered on baseband. Then, the subsequent infinite impulse-response (IIR) filter is applied and the rest of the data product algorithms can take place, with slight adjustments made to account for the higher sampling rate used. The advantage of this approach for the adjacent overloaded channels is that this approach would minimize the noise figure degradation due to adjustments to the dynamic range settings.

4.2 Measurement Tradeoffs

The addition of dynamic range adjustments allow the sensor to measure channels without triggering an overload. However, increasing the signal analyzer's attenuation and reference level settings increases the noise figure in the measured channels, decreasing the sensitivity of the sensor.

The sensor was originally designed to maximize noise performance to ensure that low level signals could be measured. However, data collected from the prototype sensors has shown that in most cases, and especially when the sensor is sited within a CBSD cluster, higher level signals overwhelm the sensor's digitizer. In these cases, the sensor should be retuned to maximize the dynamic range instead of prioritizing noise performance. Initial measurements of each sensor at a variety of dynamic range settings allow each channel to be individually tuned to the appropriate settings, minimizing the noise floor while mitigating the majority of overloaded measurements. Periodic checks will be implemented to retune these settings in the case of a change in CBSD deployments.

In addition, to further increase the effective dynamic range of the measurement campaign, several attenuation settings may be interleaved. For example, the sensor may take data products with 0 dB attenuation, then 10 dB, then 20 dB, before repeating. In this case, the overall duty cycle of measurements remains the same, but the duty cycle of measurements of the same signal analyzer settings will increase with the number of interleaved settings. However, this would provide a way to measure both high power and low power emissions without needing to frequently adjust for signal levels. The advantages and disadvantages of this strategy will be evaluated for each sensor individually based on the expected signal environment and especially in the total range of power expected.

Finally, the front-end filter was altered to help alleviate the overloads in the upper CBRS channels, but the upper channels are still measured. The effects in terms of uncertainty are being analyzed so that the appropriate qualifications can be added to this data. It should be noted here that the noise floor and calibration are deceiving. The calibration offsets follow the roll-off of the filter because they are referenced to the antenna terminal, which passes through the filter and is subject to the additional attenuation. However, the thermal noise in these channels is generated predominantly by the low noise amplifiers (LNAs) which is behind the filter and is not subject to the attenuation. The power values are still referenced to the antenna terminal despite originating on the opposite side of the filter. The result is a stepped noise floor which follows the mean attenuation of the filter. However, emissions collected through the antenna will be subject to the filter attenuation, resulting in a sloped signal to interference plus noise ratio (SINR) through the 10 MHz channel.

The dwell time for each channel was chosen to be 4 seconds to ensure capture of expected signals in the environment

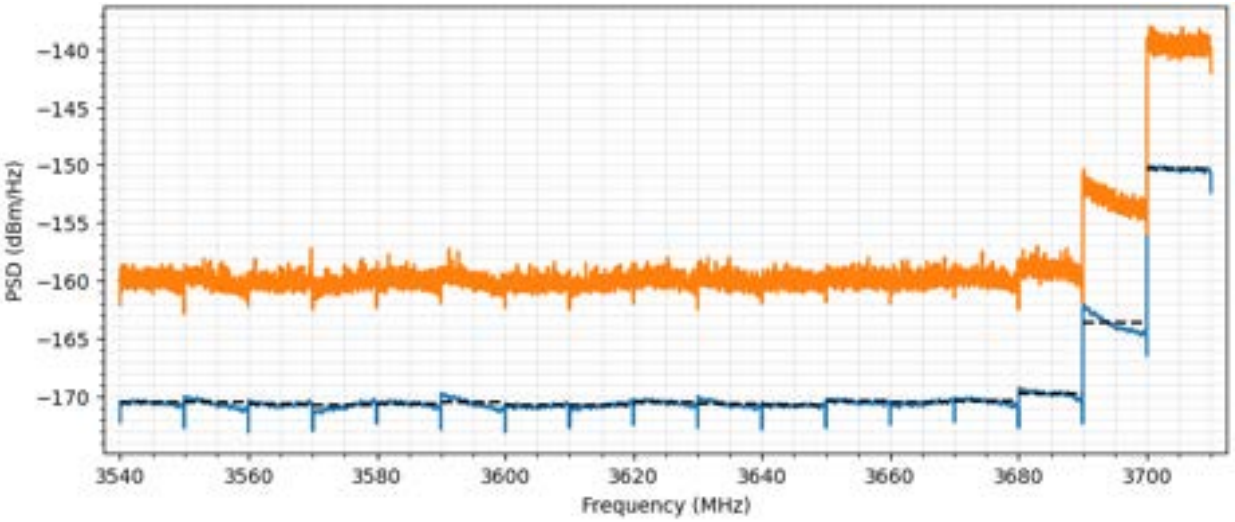

Figure 4.1: power spectral density (PSD) noise floor measurement from the sensor with the redesigned filter.

should they be present. However, this can lead to misleading overload indicators in the collected data. The signal analyzer will indicate an overload has occurred if a single sample reaches the maximum ADC value while it is digitizing. However, in observations from the fielded prototypes, it is noted that most ambient signals in the environment trend towards a distribution similar to a Rayleigh distribution, which does not have a defined peak-to-average ratio. The peak measured power is statistically defined by the dwell time and sample rate for a given acquisition. With a 56 MSps acquisition of the ADC and a 4 second acquisition time, this leads to a significantly large number of samples, over 200 Msamples. Because of this and because an overload indicator is triggered if a single one of these samples exceeds the ADC limit, overloaded data may be disregarded even if it is still significant. Users of the resultant data may wish to use the amplitude probability distribution (APD) data product to define their own "quasi-overload" indicator to reclaim data that is still significant.

4.3 Sensor Command and Control

The sensor software remains functionally the same as in the previous Test Plan [2, Chapter 5]. The spectrum-characterization and occupancy sensing (SCOS) Sensor software runs on the SPU computers and is scheduled and collected from SCOS Manager software residing on NTIA ITS servers. Sensors are provisioned using Ansible, which allows individual provisioning of sensors including calibration files and data product action parameters.

Sensor data acquisitions are locally stored until they can be retrieved by the manager software. The data and configuration files are now stored in the main volume of the processor to prevent the Docker container from exhausting its allocated resources in the case of extended network outages where data acquisitions accumulate before retrieval.

Improvements to sensor monitoring, alerting, and failure recovery have been made. To incorporate the addition of the TEC thermal management system, the additional temperature and humidity sensors have been added to the metadata as well as the current status of the TEC heater and cooler. The additional switch configurations have been added to the software to support the additional hardware functionality inside the sensor.

These metadata values have the ability to trigger alerts when ingested by SCOS manager. The appropriate alerting is being investigated and could include high temperature warnings, both in the components or in the sensors, inconsistent temperature differential between the TEC intake and outflow, or current status conditions of the processor such as high

disk utilization.

With the prototype sensors being deployed for well over a year, the failure analysis of these sensors has been included in the software to support additional automated failure recovery, including additional freeze-states and power up procedures of the signal analyzer. Throughout the deployment period of the prototype sensor, all issues have been able to be resolved remotely through the software and backhaul connection.

To support the overall data pipeline, the SCOS manager software will directly extract data from the incoming signal metadata format (SigMF) files and insert it into an Structured Query Language (SQL) database, which will allow easier querying and manipulation of the data as well as translation to a final data repository to be released.

4.4 Upgrade of Prototype Systems

The prototype sensors were originally envisioned to give an experimental look into a live CBRS environment. However, they have proven to be stable and reliable as installed. Instead of reclaiming and installing the next generation of sensor in their place, the current deployments will instead be upgraded to match the RF performance of the new sensors. To perform this upgrade the following steps will be taken. It should be noted that while many of these steps can be run remotely, all will be performed while the physical upgrade takes place in case of unexpected complications.

First, the sensor processor will be hardened to meet the requirements of the Authority to Operate (ATO) in place for the production sensors, which will involve applying the necessary Security Technical Implementation Guides (STIGs). Once this process is completed, the latest version of SCOS will be installed to match the same software revision as the new sensor iteration. Finally, the backhaul modems will be reconfigured to the production network. Once complete, the system verification will be run to ensure that the sensor is still remotely configurable.

Next, the hardware modifications will take place. The backend hardware has proven reliable, so the upgrades will match only the RF hardware. The filter in the preselector will be replaced with the new filter to help mitigate the upper channel overloads. It will also be ensured that the noise diodes in the preselector are replaced with noise diodes which have been calibrated in the National Institute of Standards and Technology (NIST) radiometer.

Once complete, the RF verification portion of the verification bench testing and acceptance (VBA) process will be rerun to calculate the gain offsets to translate the calibration plane. These offsets will then be loaded into SCOS and the appropriate validation steps will be completed.

These upgrades will proceed in stages. First the local sensor at Green Mountain Mesa (GMM) will be upgraded to ensure the process accounts for all modifications needed to bring the sensor up to production status and upgrade pitfalls are identified and mitigated. Upon completion, the sensors in the Norfolk, VA region will be upgraded, also bringing them to production status. These sensors will be registered in the production network and will begin taking production data, along with future deployed sensors. It should be noted that while the GMM sensor has served as a local sensor giving the opportunity for failure analysis where physical intervention is more convenient, it may be redeployed to a new location once it is upgraded to a production sensor.

Chapter 5

Sensor Data

The sensor data are collected in the same manner as in the previous NASCTN CBRS SEA Test Plan for Always-on DPAs [2]. The sensor command and control approach is described in detail in the previous Test Plan, and updates to it are provided above in Section 4.3. The edge compute algorithm executed by the SCOS Sensor software produces a diverse set of time- and frequency-domain power measurements and statistics from its input, a time series of IQ samples recorded by the signal analyzer. The resulting processed data and accompanying metadata are stored as SigMF archives, and raw IQ recordings are discarded. Most technical details of the edge compute algorithm are unchanged from their descriptions in the previous Test Plan [2]. Updates to this algorithm and its outputs are summarized in Section 5.1. Each of these is specifically justified below, but in general, the changes are results of optimization and responses to analysis needs over the course of prototype sensor deployment. Section 5.2 provides brief summarizing descriptions of the primary data payload components and updated example plots of each. Finally, Section 5.3 describes the approach to data ingestion and aggregation into a central database, as well as potential avenues for aggregate analyses.

5.1 Changes to Sensor Data Acquisition

Here we summarize changes made to the data acquisition routine, compared to the detailed description available in the NASCTN CBRS SEA Test Plan for Always-on DPAs [2].

- **New channel power summary statistics** have been added to each measurement's metadata. Median and median-of-mean summary statistics per channel are now recorded alongside the existing RMS and peak statistics. The median statistic is computed across instantaneous power samples for each acquisition, while the other three statistics are computed from the results of the power-vs-time mean and maximum detectors which are components of the data payload. The addition of these traces create new opportunities for lightweight longitudinal and aggregate analyses.
- **New PSD traces** have been added to supplement the existing mean and maximum traces. Observations of APDs recorded by prototype sensors indicate a tendency towards measuring Rayleigh-type distributed powers. This type of distribution has no theoretical maximum, and the peak-to-average ratio is statistically defined by the time-duration of the capture and the sampling rate. Because of this, and the relatively long capture durations, the maximum trace can be deceptively high depending on the type of analysis performed. Therefore, traces indicating the following percentiles for each fast Fourier transform (FFT) bin were added: 25^{th}, 50^{th} (median), 75^{th}, 90^{th}, 95^{th}, 99^{th}, 99.9^{th}, and 99.99^{th}. These new traces are shown below in Figure 5.4. Note the large gap between the maximum and 99.99% traces.
- **PSD frequency bin size has been increased from 16 kHz to 80 kHz** to keep overall data sizes small while accommodating the new PSD traces. This choice retains frequency resolution sufficient to resolve 3rd-generation partnership project (3GPP) guard bands while still spanning each 10 MHz channel with an integer number of bins, to support data aggregation.
- **Total measurement span has been expanded** to include additional channels adjacent to the CBRS band. The measurement band now spans 3530–3710 MHz. Although the sensor design was optimized for high-sensitivity measurements within the CBRS band, the sensor is still able to take calibrated measurements across this extended

frequency range. Due to measuring these additional channels, the sweep time (i.e., channel revisit time) has increased to approximately 100 seconds.

As stated above, additional dynamic range adjustments can be made channel-by-channel to accommodate emissions of a wide range of powers. These settings are configured as parameters of the SCOS action configurations and will be maintained for as long as the action is scheduled, but can be updated at any time through rescheduling the software if needed. The appropriate calibration values for the given settings are also packaged with the metadata. Based on laboratory data, it was found that a ratio of 25 dB between the signal analyzer's reference level and attenuation settings gave consistent dynamic range across the different settings, thus this ratio is maintained for all attenuation levels used in the actions.

5.2 Summary of SEA Edge-Compute Data Processing and Products

The edge-computed data product consists of the same four categorical data products as described at length in the previous Test Plan. All are derived from 4 second long stepped-frequency IQ recordings in 10 MHz channels spanning the CBRS band, as well as two 10 MHz channels below and one 10 MHz channel above the CBRS band. Thus the frequency band measured by each sensor is 3530—3710 MHz. Each 10 MHz capture is sampled at 14 MSps, then filtered to 10 MHz with the same IIR filter detailed in the previous Test Plan before individual data products are computed (note that the initial downsampling from 56 MSps to 14 MSps is performed within the signal analyzer before passing the IQ stream to the processor).

The channels are sampled sequentially, from low to high frequency. Data processing occurs concurrently with IQ acquisition to reduce the overall measurement duty cycle, i.e., channel N data is processed while channel $N + 1$ is sampled. With the eighteen 10 MHz channels being measured, the approximate time per acquisition is 100 seconds.

5.2.1 Power Versus Time

Unchanged from its description in the previous NASCTN CBRS SEA Always-On DPA Test Plan [2], the power versus time (PvT) data product indicates the temporaly varying power across the entire 4 second capture at 10 ms resolution. Individual IQ data points are calculated as instantaneous power with a 50 ohm impedance of the signal analyzer, then measured using both peak and RMS detectors over 10 ms intervals through the 4 second capture, based on the frame length of long-term evolution (LTE). This is supplemented by the PFP as described below which gives insight on a shorter time scale.

5.2.2 Periodic Frame Power

The periodic frame power was devised as a data-compressive approach to analyze channel occupancy on timescales smaller than the 10 ms provided by the PvT measurements. It is calculated from time-series power data computed from the IQ recording with a bin size equivalent to the shortest anticipated occupancy event- in this case, $\frac{1}{56}$ ms, the shortest transmission duration allowed by Fifth Generation of cellular mobile communications (5G) standards. For this stage, RMS and peak power detectors are used. These power series are segmented into 10 ms reference frames, and {minimum, RMS, maximum} power detectors are applied across samples sharing an index within these frames. The result is a set of 6 traces of cyclostationary power for each measurement of a 10 MHz channel.

The periodic frame power has proven invaluable in analyzing emissions with cyclostationary occupancy characteristics, including TDD cellular downlink and Radar 1 signal types. It is retained without modification from its description in the previous NASCTN CBRS SEA Always-On DPA Test Plan [2]. Periodic frame power traces can be displayed as a cyclostationary persistence plot, showing cyclical occupancy signatures captured in a single measurement of a single channel. Aggregate analysis can use correlation to align consecutive captures and assess network loading over larger

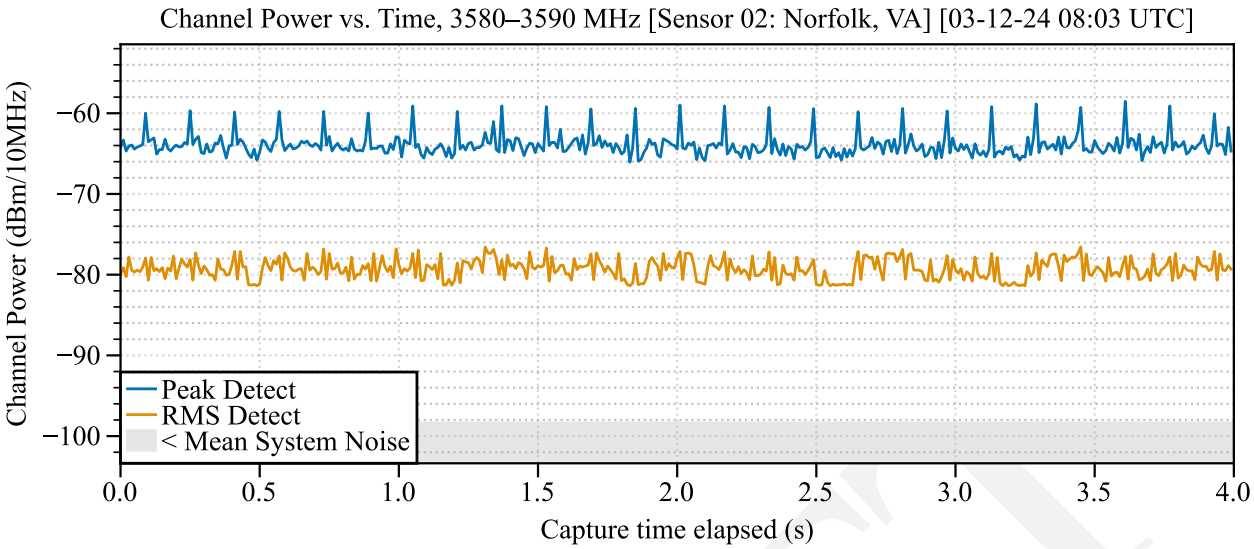

Figure 5.1: Example power versus time data for the 3580-3590 MHz channel, occupied by TDD downlink transmissions, measured by a prototype sensor in Norfolk, Virginia.

timespans.

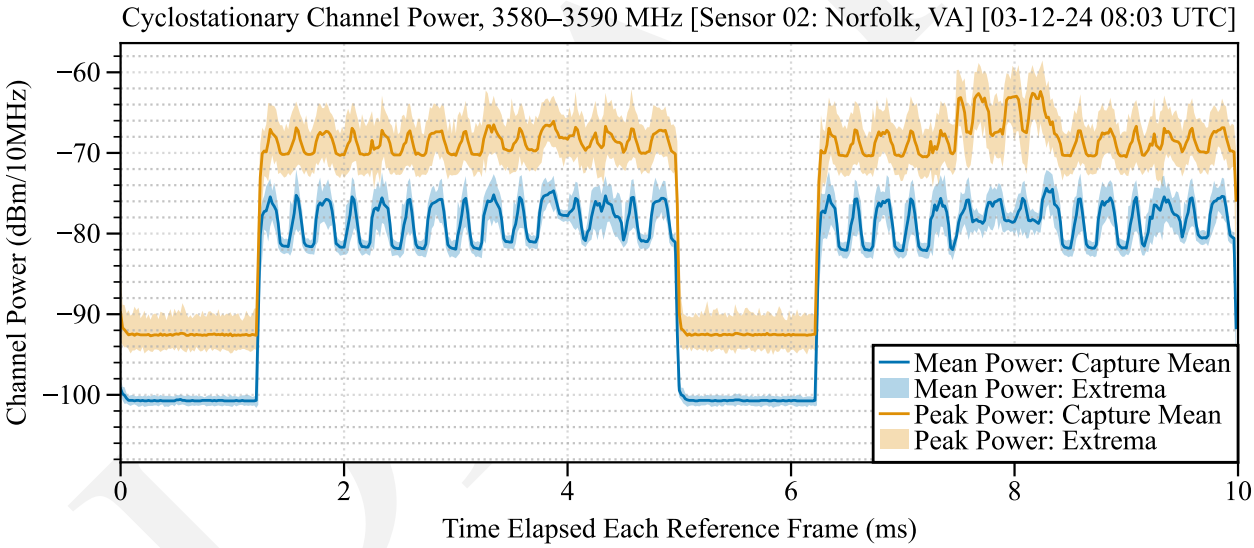

Figure 5.2: Example cyclostationary channel power data for the 3580-3590 MHz channel, occupied by TDD downlink transmissions, measured by a prototype sensor in Norfolk, Virginia.

5.2.3 Amplitude Probability Distribution

The APD is a complementary cumulative distribution function (CCDF) of the instantaneous power samples in the 10 MHz channel throughout the 4 second capture. The APD can be used to easily compare the distribution of measured power to known distributions, or to other measured signals. Since a traditional estimate of the APD would result in an array of values equal to the original IQ length, it is instead calculated by the edge compute algorithm in 1 dB amplitude bins. This accomplishes necessary reductions in total algorithm computation time and data product file size. This effectively reduces the overall resolution of the APD, but was deemed an important tradeoff to process the

data in a timely fashion. The resolution of a full APD would likely be excessively great for any analysis on the final data. Additionally, the large data sizes associated with it would make any aggregate or longitudinal analysis more difficult.

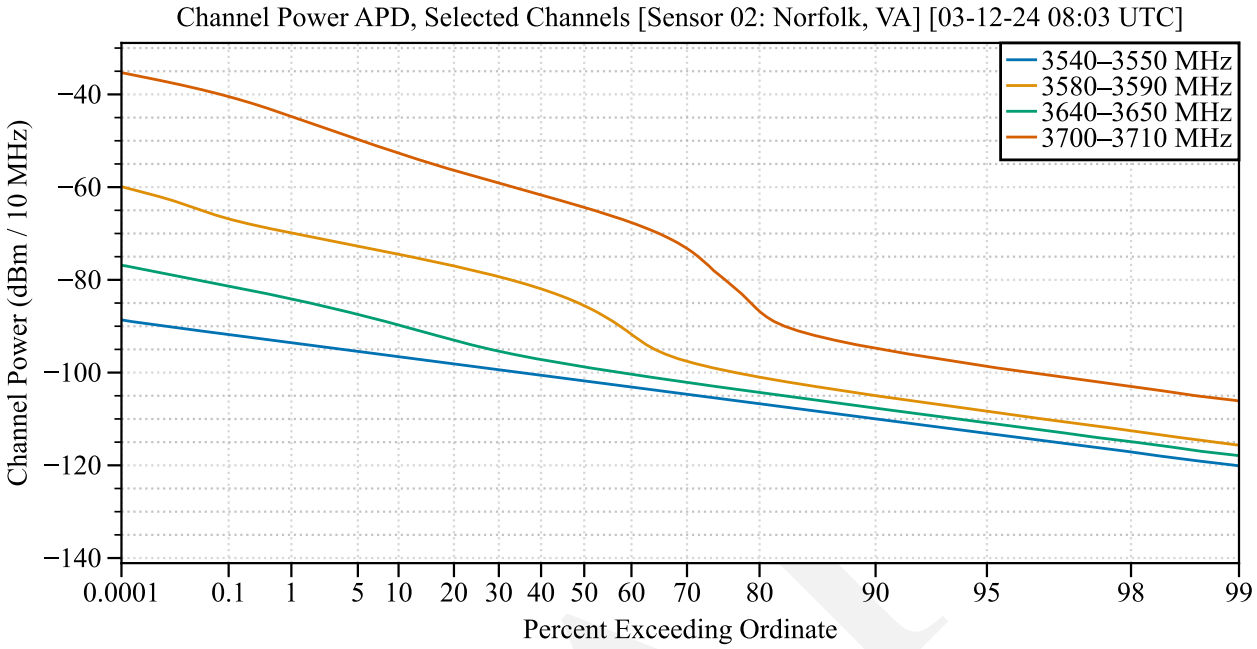

Figure 5.3: Example APDs for selected channels exhibiting different occupancy characteristics, measured by a prototype sensor in Norfolk, Virginia.

5.2.4 Power Spectral Density

The power spectral density measures the distribution of RF power present in the recorded signal as a function of frequency. In our approach, a desired frequency resolution of 80 kHz is achieved by segmenting the recorded IQ waveform into 175-sample segments before computing the discrete Fourier transform (DFT) of each segment. This approach is derived from the Welch method, but computes additional statistics beyond the mean. The reduced frequency resolution accommodates the data size increase due to the added statistics across the DFTs. The added traces provide insights into the occupancy characteristics of measured signals, as functions of frequency. This essentially bridges the gap between the previous PSD implementation and the APD. The sacrificed frequency resolution does not obscure any signal features which were of interest for analysis. We also note that computationally-expensive resampling was intentionally avoided, which could have allowed for the frequency binning to align with expected signal structures.

5.3 Data Aggregation

As described above, the data products are packaged into SigMF files on the sensor and are retrieved by the manager software, which extracts the data and inserts it into an SQL database. This database allows for data to be queried in a simpler fashion and for easier bulk data manipulation. This database was designed to match the data product in MySQL allowing a one-to-one match between the incoming meta-/data to the database.

The data in this database will be analyzed and validated to ensure validity. Any sensitive data will be stripped from the database and any data products can be reshaped as needed. Further analysis on these data products may be conducted resulting in a final dataset to be published. This final dataset will then be migrated to a public accessible database

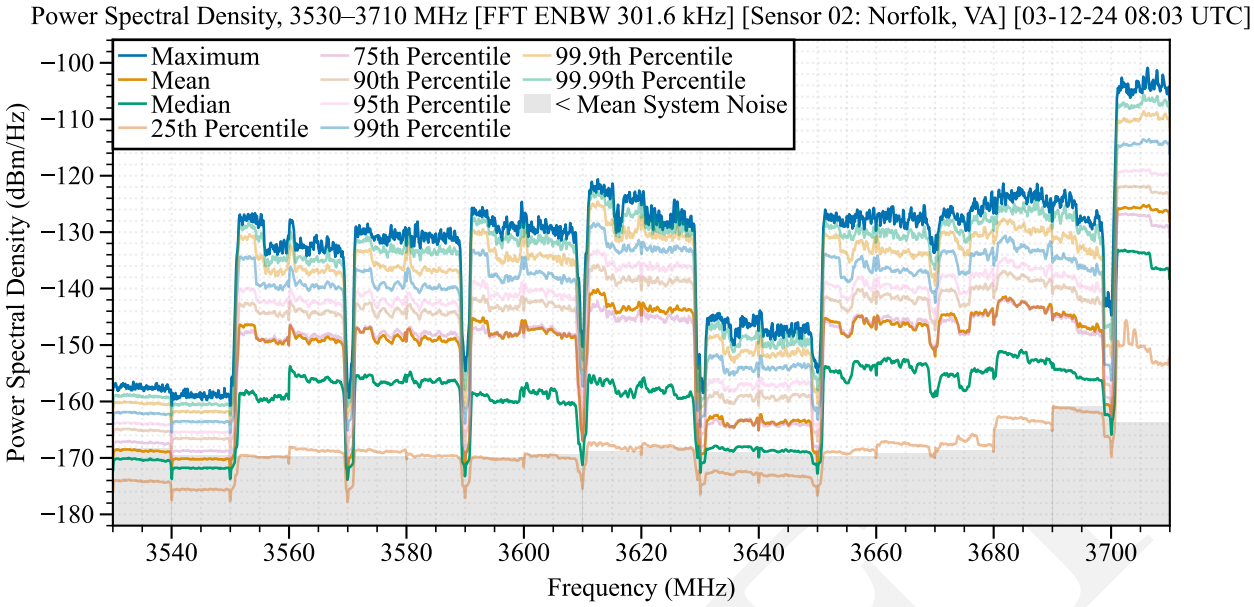

Figure 5.4: Example PSD data product produced by a prototype sensor in Norfolk, Virginia.

at the conclusion of the project to be maintained on a NIST-administered Amazon Web Services (AWS) Relational Database Service in accordance with the NIST Open Data Program and Data Retention Guidelines.

5.3.1 Short-Term Dynamics

The collected data products were designed to give insights to both short-term dynamics of the CBRS ecosystem as well long term trends. Towards short term dynamics, short sequential data collects can shed light on transients in system, such as rapid reductions in emissions, on the order of acquisition duty cycle (i.e. 100 seconds). These events can then be strung together to create a "day's view" of the CBRS ecosystem as observed by the sensor as can be seen in the Figure 5.5 below.

5.3.2 Long-Term Trends

In addition to short term dynamics, the data can be aggregated over longer time scales to give insights to the long-term behavior of the ecosystem. In the simplest form, a detector (mean or peak) can be applied across the PvT for an entire day and investigated over a period of months to give insights to trends that occur on longer time scales, e.g. cyclical across a week-long period. Other analyses may include investigation of the relation between mean and peak powers in individual data collects aggregated over long periods of time to isolate behaviors in the system such as in Figure 5.6 below. Note that an anomaly is easily detected in these data, corresponding to a testing event observed by the sensor.

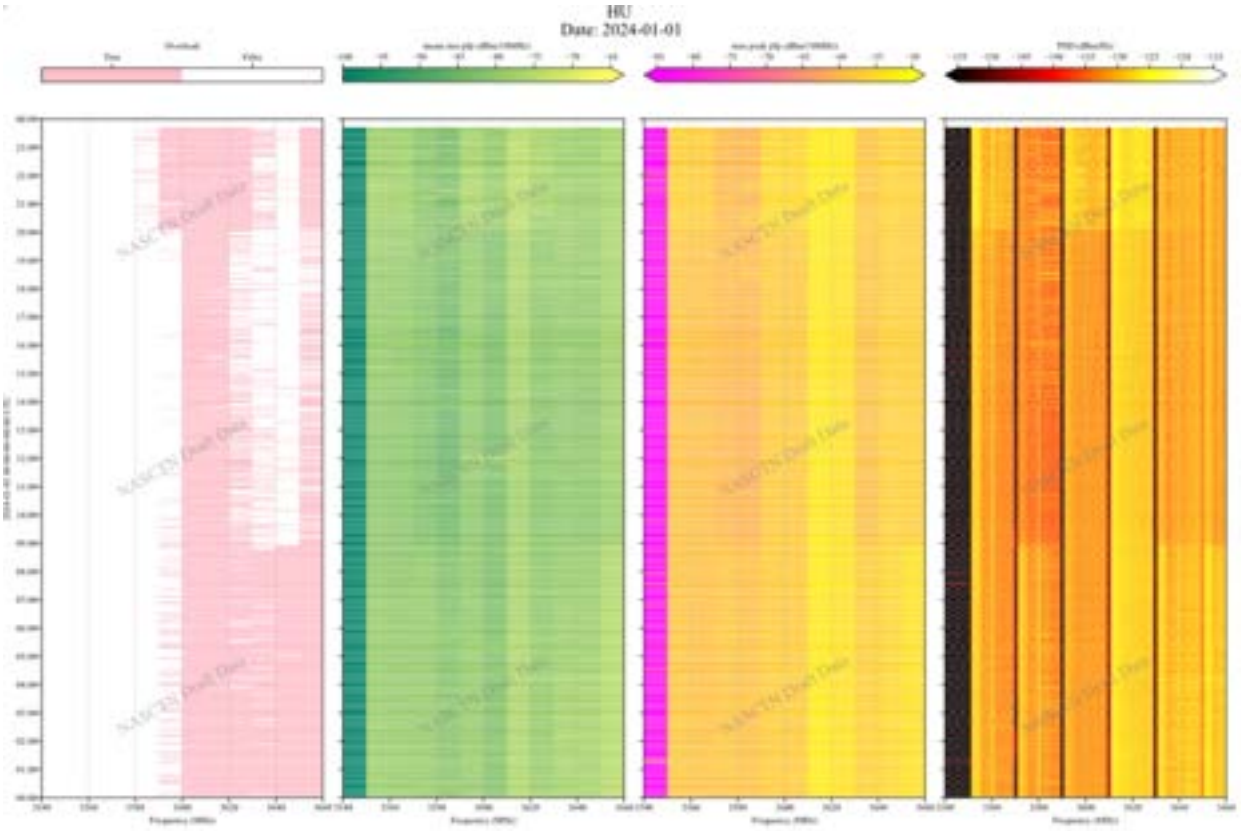

Figure 5.5: Example day cycle trend plot derived from PFP and PSD data products produced by a prototype sensor in Norfolk, Virginia.

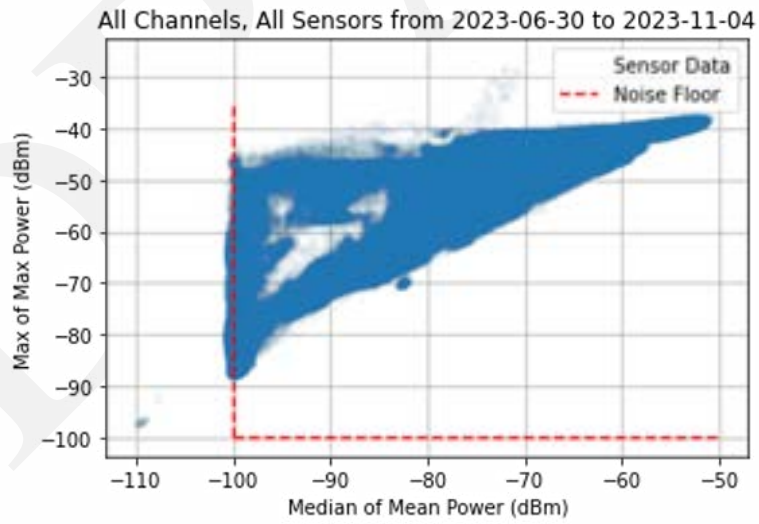

Figure 5.6: Example longitudinal max-of-max versus median-of-mean plot derived from PFP data products acquired on all three prototype sensors across approximately 6 months.

Chapter 6

Metrology Sensor Demonstration

In October of 2022, two prototype sensors were deployed in the Norfolk, VA area and one prototype sensor was deployed on Green Mountain Mesa (GMM) in Colorado. Additionally, after the sensor design revision documented in Chapte 4, a fully outdoor sensor was also deployed to GMM for failure effects analysis and test the new design. These sensors have served as exploratory looks into the Citizens Broadband Radio Service (CBRS) ecosystem and have served to refine and experiment with the data products as well as provide installation experience and failure mode analysis for the sensor architecture and software. In the period since installation, these sensors have not required physical intervention with all sensor issues being solved remotely. In addition, they have continually collected data in a development environment throughout which has proved invaluable to understanding the ecosystem and possible analysis techniques.

6.1 Sensor Installations

Three prototype sensors were deployed with the intent of testing the sensor architecture and gaining insight to the CBRS ecosystem. Two were deployed near the East-1 Dynamic Protection Area (DPA) in Norfolk, VA to observe a dynamic environment and test remote administration while one was deployed locally at GMM to observe a static environment with the ability to manually intervene if necessary. The installation details are described in the previous test plan[2]. The data taken from these sensors has shaped much of the analysis and refinements to the data products as described below.

Based on failure mode analysis, observations of the environment, and evolving deployment requirements, the sensor design was refined into the production sensor described in Chapter 4. Upon construction of these sensors, one was deployed on GMM in January 2024 with the similar intent as with the prototypes, namely test the new sensor design in a manner where physical intervention was feasible if needed. Of particular interest was the effectiveness of the all-outdoor design of the new sensor needed to expand the possible deployment options which successfully kept the components within the signal processing unit (SPU) at their operating temperatures in an ambient condition of -26° C.

Deployment of the new sensor in a live CBRS environment has also given the ability to test the radio frequency (RF) path upgrades, namely the new filter design meant to alleviate high power emissions in the channel directly above CBRS. Figure 6.2 below showcases the power spectral density (PSD) plot from the production sensor installed on GMM, indicating that the redesigned filter is successfully mitigating overloaded conditions in the upper CBRS channels. Note that the noise floor, as indicated by the dashed line, in the upper-most channels increases and the PSD trend slopes with the filter rolloff. This and dynamic range testing on the prototype sensors have formed the basis for the strategy to effectively measure all CBRS channels despite the power levels incident on each sensor.

The deployment of these sensors has formed a pipeline of software evolution to bring new features and improvements successfully into the sensor architecture. The three sensors at GMM provide a gradual step towards production with software improvements developed on the broken out revision 0 system, then transitioned to the prototype sensor for initial testing. Once these changes prove stable, they are pushed to the Virginia sensors to collect data in a live,

Figure 6.1: Installation of the new sensor design at GMM with the SPU (left) and preselector (right). Note the preselectors for other sensor revisions are located near the sensor revision 1 preselector as well.

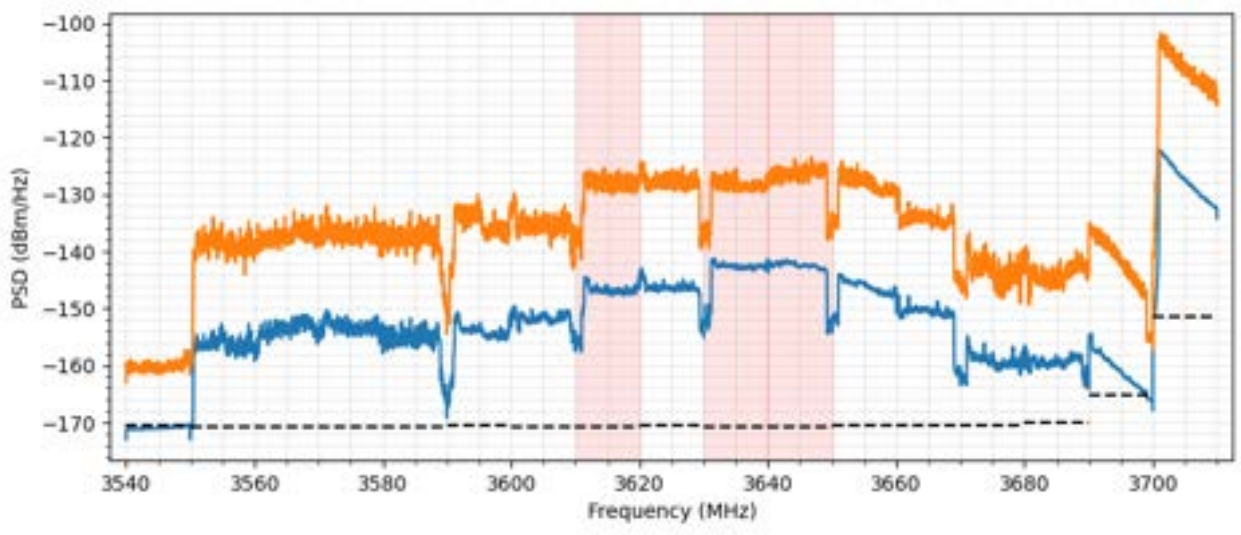

Figure 6.2: Example PSD data taken by the GMM sensor revision 1.

dynamic environment anticipated for future deployments. In addition, the changes are also pushed to the production sensor at GMM to ensure that the software is compatible across the sensor revisions and is hardware agnostic.

Given the success of the prototype sensors, they will be upgraded to match the RF and security requirements of the production sensors as described in Chapter 4 instead of redeploying a new sensor to the same location, effectively increasing the number of production sensors available. As they will still require a split-deployment, locations will still be sited that support this. In addition, as the software reaches full maturity and development transitions to maintenance,

the sensors currently deployed at GMM will be re-sited as new deployment locations are secured.

6.2 Observations from Deployed Sensors

With well over a year's worth of data collected from the fielded sensors as well as data collected from site surveys of potential deployment locations, exploratory analysis techniques have begun to form to showcase the usability of the dataset in analysis into the effectiveness of the CBRS ecosystem. Several such techniques are showcased here. Note that these analyses were formed throughout the evolution of the software and data products, so the data showcased may have different shapes dependent upon when it was taken.

6.2.1 Day Cycle Trends

Originally shown in the first Test Plan[2], the data collected by an individual sensor can be stitched into a daily waterfall plot showing statistics across the day to the time resolution of the data collections. This analysis has proven invaluable in checking the daily trends of the individual sensors including identifying short-term changes in power in individual channels. The statistics used in these plots may be changed dependent upon the desired analysis being performed.

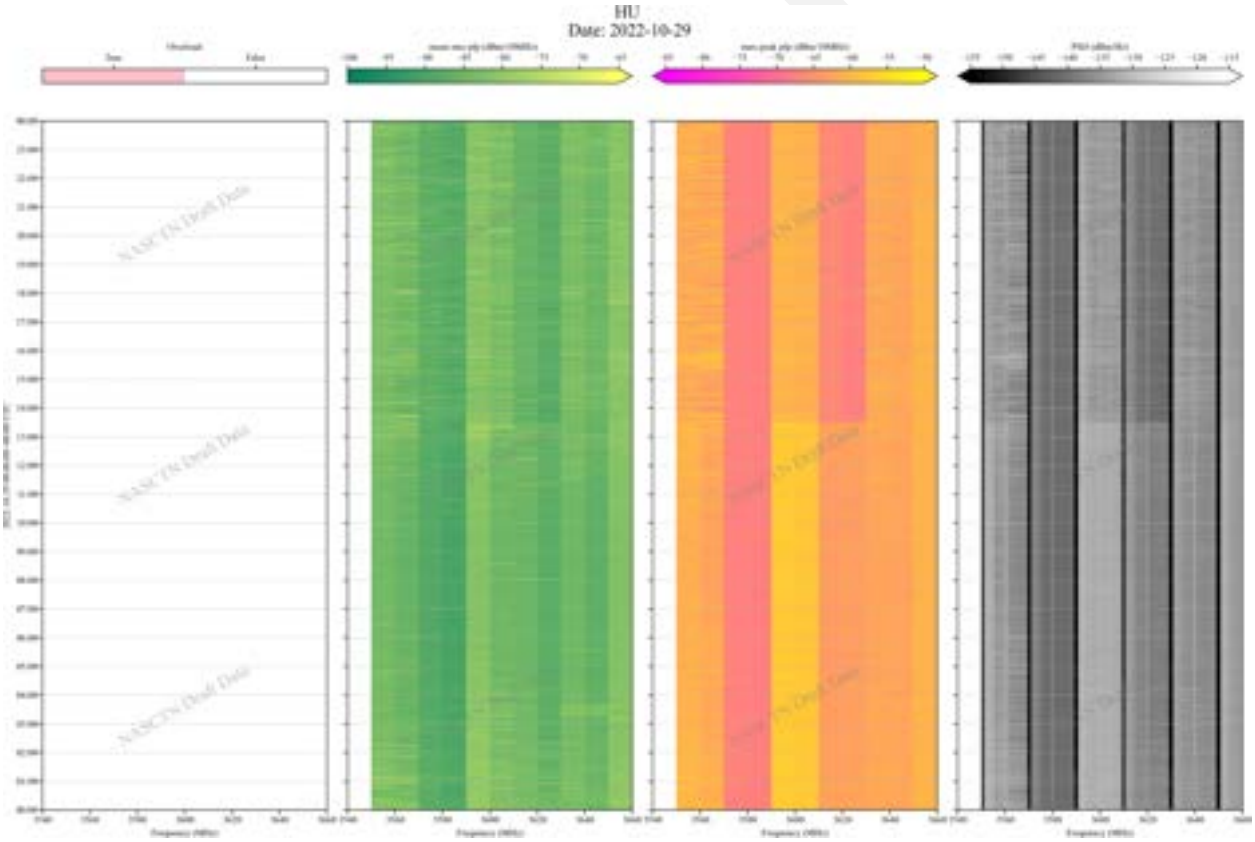

Figure 6.3: Day cycle trend plot derived from periodic frame power (PFP) and PSD data products produced by the Hampton University in Norfolk, Virginia.

6.2.2 Longitudinal Power Statistics

In addition to shorter term analyses on singles sensors, data from multiple sensors across extended periods of time can be combined to form a more longitudinal analysis. For example, Figure 6.4 below shows data from all three deployed prototype sensors for an approximately six month period.

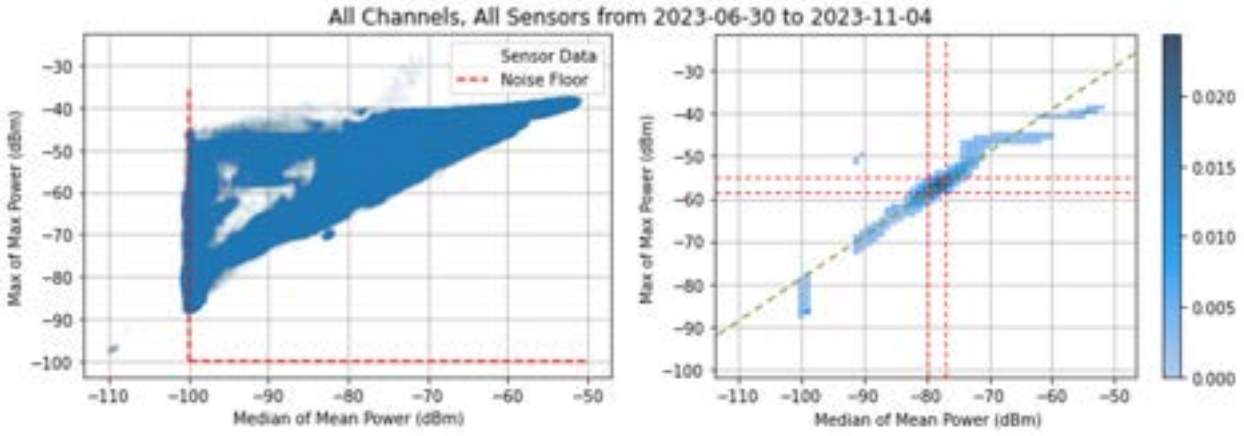

Figure 6.4: Plot of max peak PFP measurements versus median root mean square (RMS) PFP measurements across all prototype sensor for an approximately six month period. On the left shows all data while the right shows the most probable subset of data.

These plots showcase the max-of-max power versus the median-of-mean power derived from PFP measurements for each data collection across all sensors for approximately six months, effectively giving a long-term measure of the Peak-to-Average Power Ratio (PAPR). The dashed lines in the left-most plot indicate the approximate sensor noise floor, which is the lower bound for the median-of-mean axis. Due to a sensor error, uncalibrated measurements were taken immediately after a spectrum-characterization and occupancy sensing (SCOS) schedule was executed leading to the data points in the lower left, but these are easily discernible in the plot (this issue has since been resolved). The upper bound of the "triangle" is limited by the compression level of the sensors, which cannot measure above a certain power level for a given dynamic range setting. With the addition of the overload flag in the metadata, these points can be ignored, marked, or otherwise removed from the dataset. Note that the dynamic range settings were briefly adjusted during this measurement period, leading to the lightly colored region beyond this boundary, indicating the effectiveness of adjusting the dynamic range settings in mitigating sensor overloads.

This data was then further reduced to the plot on the right, showcasing the most probable data points for the sensor measurements. As each sensor is a different distance from the nearest neighbor Citizens Broadband Radio Service device (CBSD) in each channel, the data creates a trend in effective PAPR across all sensors and all channels instead of a singular point. However, this trend is relatively constant, pointing towards a generally constant PAPR across the measured CBSDs.

6.2.3 Aggregate Measurements

The placement of the three original prototype sensors resulted in measurements of predominately a single CBSD, the nearest neighbor, due to the close proximity of the surrounding CBSD cluster and the quicker rolloff of the path loss curve. However, placement of sensors further away from the CBSD cluster can form more of an aggregate measurement where there are power contributions from multiple CBSDs to the measured power statistics. This can be seen in the data taken during a site survey at Catalina Island.

In the PFP plots on the right of Figure 6.5, apparent contributions from multiple CBSDs can be seen as indicated by

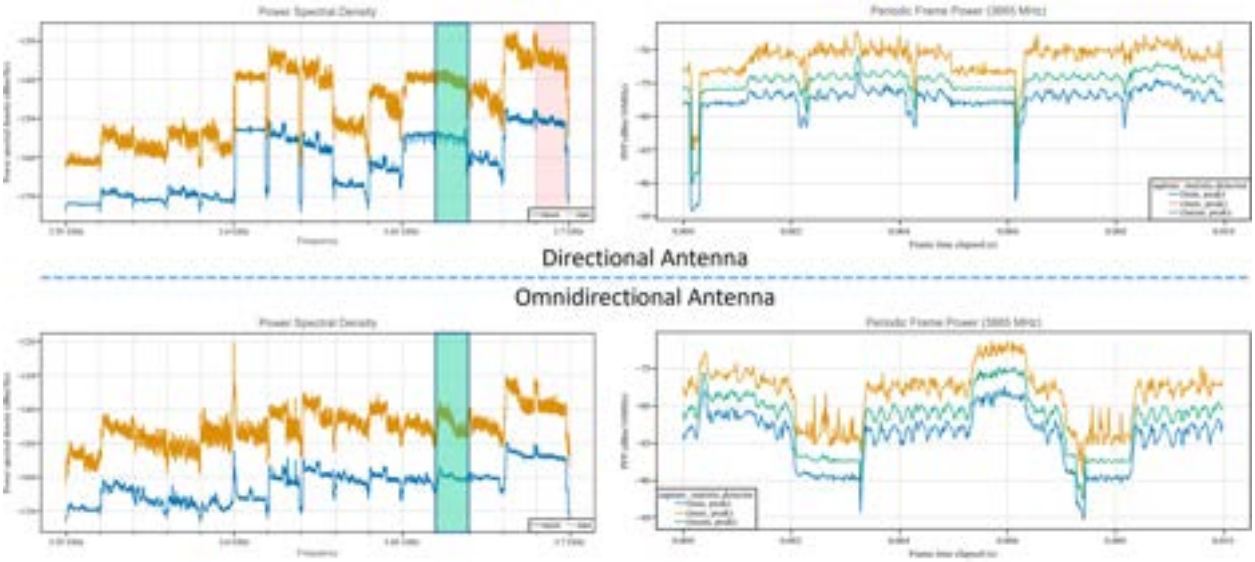

Figure 6.5: Plot of PSD (left) and PFP (right) data for the highlighted channel with both a directional (upper) and omni-directional (lower) antenna measured at Catalina Island.

the power levels exceeding the sensor noise floor in the uplink portions of the time division duplex (TDD) structure. The aggregation of the signals will be strongly dependent on the antenna pattern of the sensor and the distribution of CBSDs as demonstrated in Figure 6.6. This is apparent in Figure 6.5 as the difference in PFP structure between the directional and omni-directional antenna measurements. In addition, the measured power using the omni-directional antenna is higher in some channels than the directional antenna, indicating different contributions from the measured CBSD cluster. In the case of the Catalina Island sensor, this introduces the idea of using multiple antennas, one an omni-directional and one directional with a narrower beam-width, to help understand how these signals aggregate together. Aggregate measurements of the same CBSD cluster through different antenna patterns will help showcase how these signals aggregate over long path distances.

The aggregate measurements can help the community understand how these signals aggregate at distance when the measurement point is not in the nearest-neighbor regime. In addition to the PFP Figure 6.5 above, the amplitude probability distribution (APD) and additional percentiles in the PSD data will give insights into how the distribution of powers change as the measured emissions transition from dominated by a single CBSD to contributions from multiple CBSDs.

6.2.4 Additional Percentile Views

Based on initial collected data, the distribution of powers seen trended towards distributions which do not appear to have a defined maximum, instead being statistically determined by the capture length and sample rate. Because of this, the choice of sensor capture duration, 4 seconds, can create a deceivingly high maximum power, both in the APD and PSD data.

Instead of relying on a maximum value, which will be dependent on the measurement parameters, it may be more advantageous to use a defined percentile, which appears much more constant in the measured data and is independent of the measurement duration in distributions such as the Rayleigh distribution. This is captured across the entire 10 MHz channel in the APD data, but would also be of use at a finer spectral resolution. As such, the additional percentiles were also added to the PSD data product in addition to the traditional mean, median, maximum, minimum

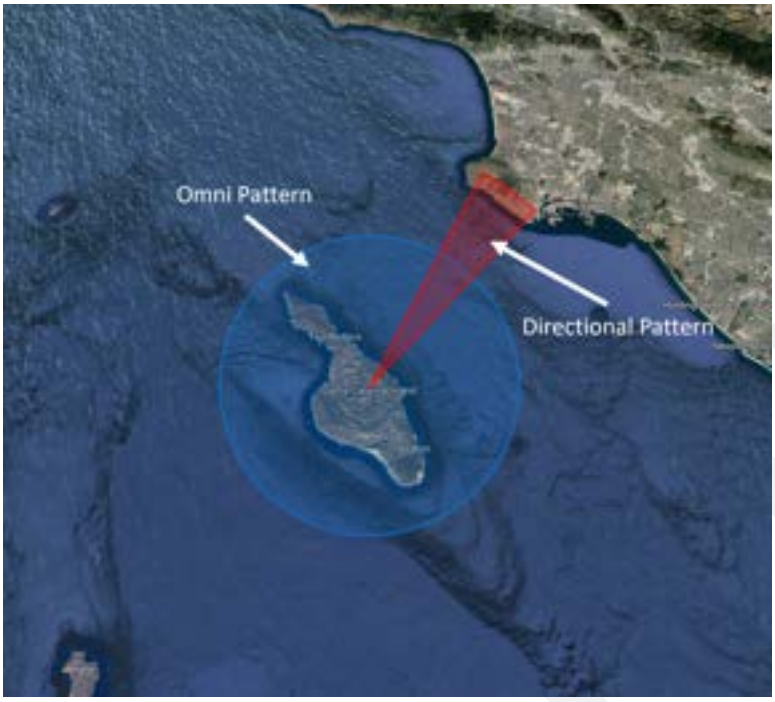

Figure 6.6: Example omni-directional and directional antenna patterns with relative gain located on Catalina Island.

(M4) trends to facilitate more insights into the power distribution between the mean and peak traces.

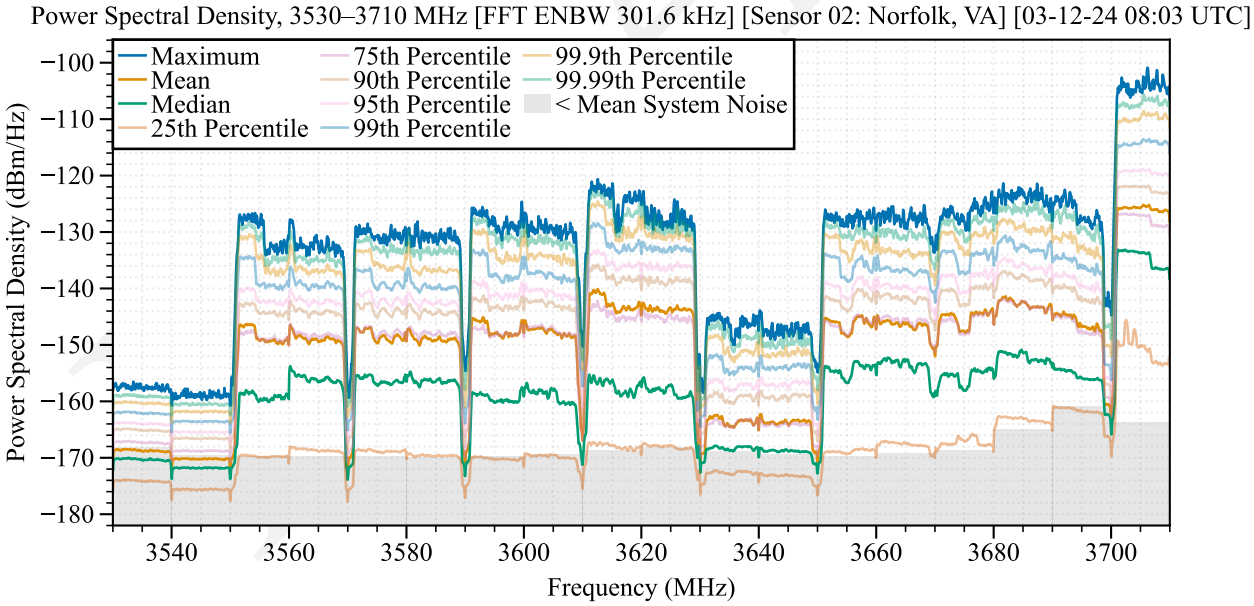

Figure 6.7: Example PSD data product produced by a prototype sensor in Norfolk, Virginia.

The additional percentiles traces fill in the gap between the mean and peak traces, showing the distribution of powers at a finer spectral resolution than the APD alone. While these do not include the full range of percentiles, discrete percentiles of interest were included (documented in Section 5.2.4). In Figure 6.7, it can be seen that there is still a

discernible delta between even the 99.99th percentile trace and the maximum trace. It is worth noting that because the PSD traces reduce the sample size per discrete Fourier transform (DFT) bin, the PAPR will differ between the APD across the entire 10 MHz band and the PSD bin width (80 kHz), supporting the need of additional percentiles between the mean and the maximum. In addition, the median and 25th percentile traces may fall below the mean system noise floor as these percentiles may fall at power levels below the Rayleigh mean of the thermal noise depending on the measured distribution of powers in a given channel (for example in the two channels below the CBRS band in Figure 6.7).

6.2.5 Deployed CBRS Technologies

The measured data can also give additional insights to the deployed CBRS technologies and trends. For example, measurements taken in the McLean, VA region tend to show 20 MHz allocations based on the location of guard bands. On the other hand, measurements along the West coast tend to show predominantly 10 MHz allocations, again based on the location of guard bands.

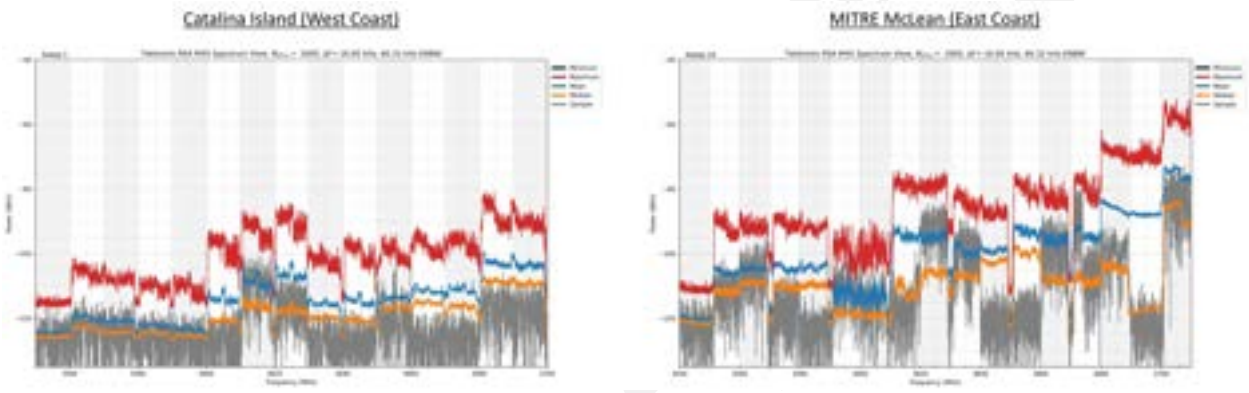

Figure 6.8: PSD data on the West and East coasts showing different allocation trends.

In addition, the data can showcase the utilization of the CBRS band as it is affected by the presence of nearby ground-based DPAs (GB-DPAs). As the lower channels have emission limits, the CBSDs allocated in these channels necessarily must be either further away or reduced in power to meet these limits. In contrast, the upper General Authorized Access (GAA) channels which do not have the same emission limits can be much more utilized near the GB-DPAs. This can be seen in the data from both Camp Pendleton and Miramar in Figure 6.9.

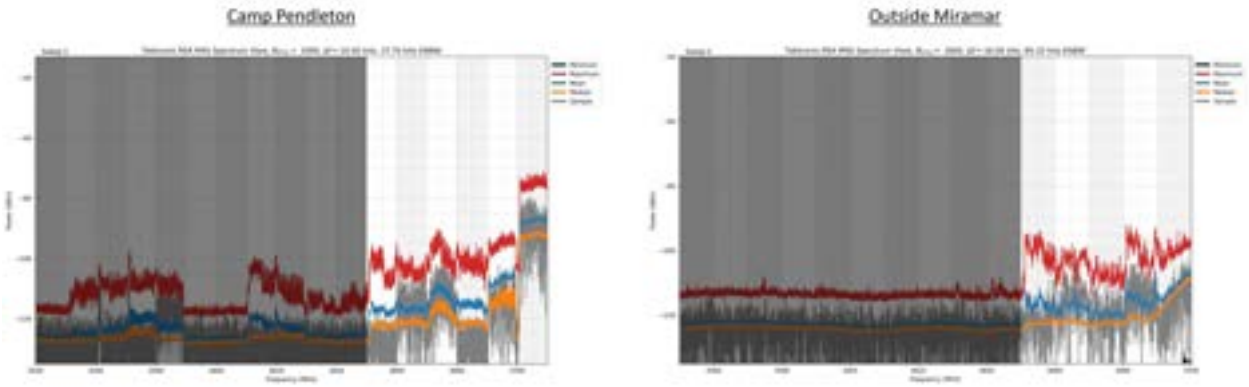

Figure 6.9: PSD data within or near GB-DPAs along the West coast.

While a more in depth signal classification analysis would be required to confirm, the PSD plots do give insight to the frequency structure of the signals, possibly identifying different types of deployed technologies. For example, in Figure 6.10, the PSD of the lower 40 MHz appears different from the typical structure seen in other allocations as evident by the apparent presence of pilot tones throughout the 40 MHz allocation. However, in the third channel, the frequency structure appears more like the 3rd-generation partnership project (3GPP) signal structure seen throughout the rest of the band. Note that this may also be analysed as a possible aggregation of different technologies in the CBRS band.

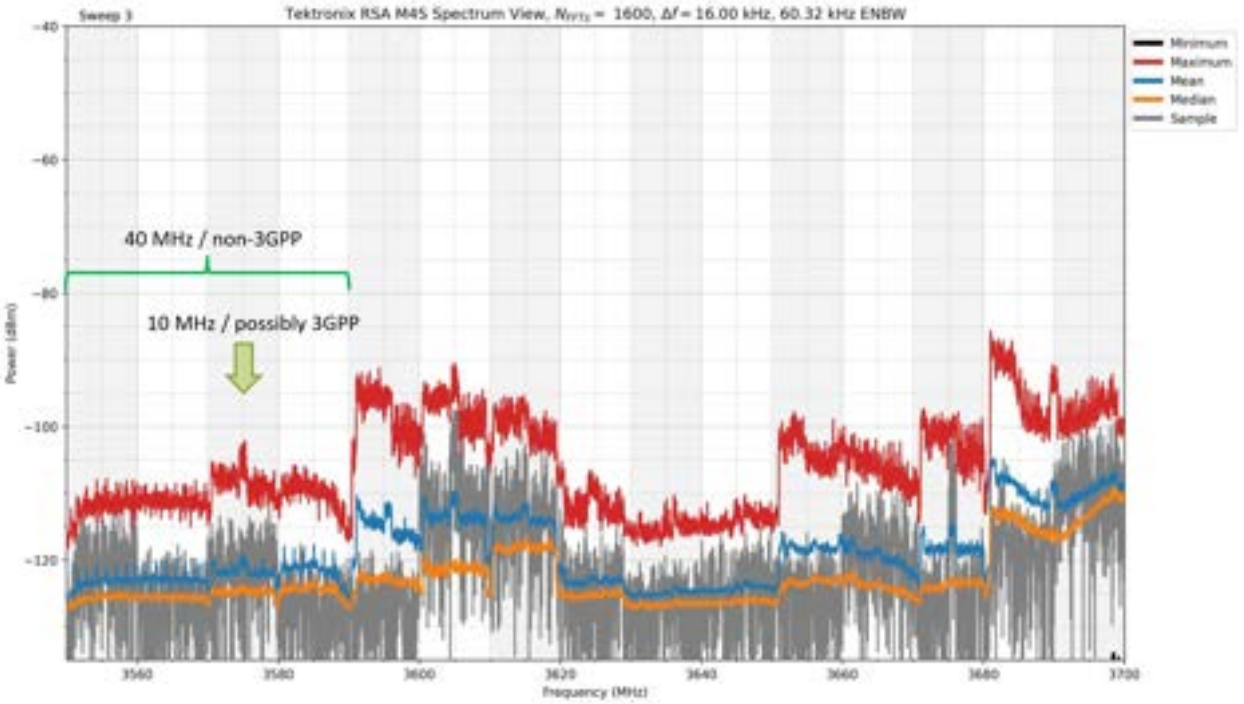

Figure 6.10: PSD data taken along the West coast showing different frequency structure in the lower CBRS bands.

6.2.6 In-Depth Technology Analysis

More in-depth analyses can also be performed on the data. For example, the non-shared channel resources can be mapped onto the PFP data for 3GPP structures to give insights to how the channel is performing. Based on the observed PFP data, it is likely that many of the 3GPP deployments are following TDD configuration 2, special subframe configuration 7. The resource grid for this configuration can be seen in Figure 6.11 (reference: dhagle.in/LTE) for a 10 Mhz channel. The shared channel (PSSCH) contains the downlink resources used for data transmission to user equipment (UE), which may be used depending on the network activity in the channel. However, the non-shared channel resources (PSCCH, PBCH, etc.) may be used irrespective of downlink activity, leading to a different relative power for resource elements containing non-shared channel channel resources.

The resource grid in Figure 6.11 can be used to create a symbol-by-symbol mask of the non-shared channel resources in the time domain. This mask can then be aligned to the PFP data product for the respective 3GPP allocations as can be seen in Figure 6.12. There is good correlation between the mask and PFP traces, indicating that the symbols containing non-shared channel resources are either used more frequently than the purely shared channel symbols or are transmitted at a higher power level. In addition, the relative delta in power between the masked and non-masked symbols may give insight to the current shared channel utilization in the given time period (i.e. if the shared channel is highly utilized, then the delta between the symbols would be less).

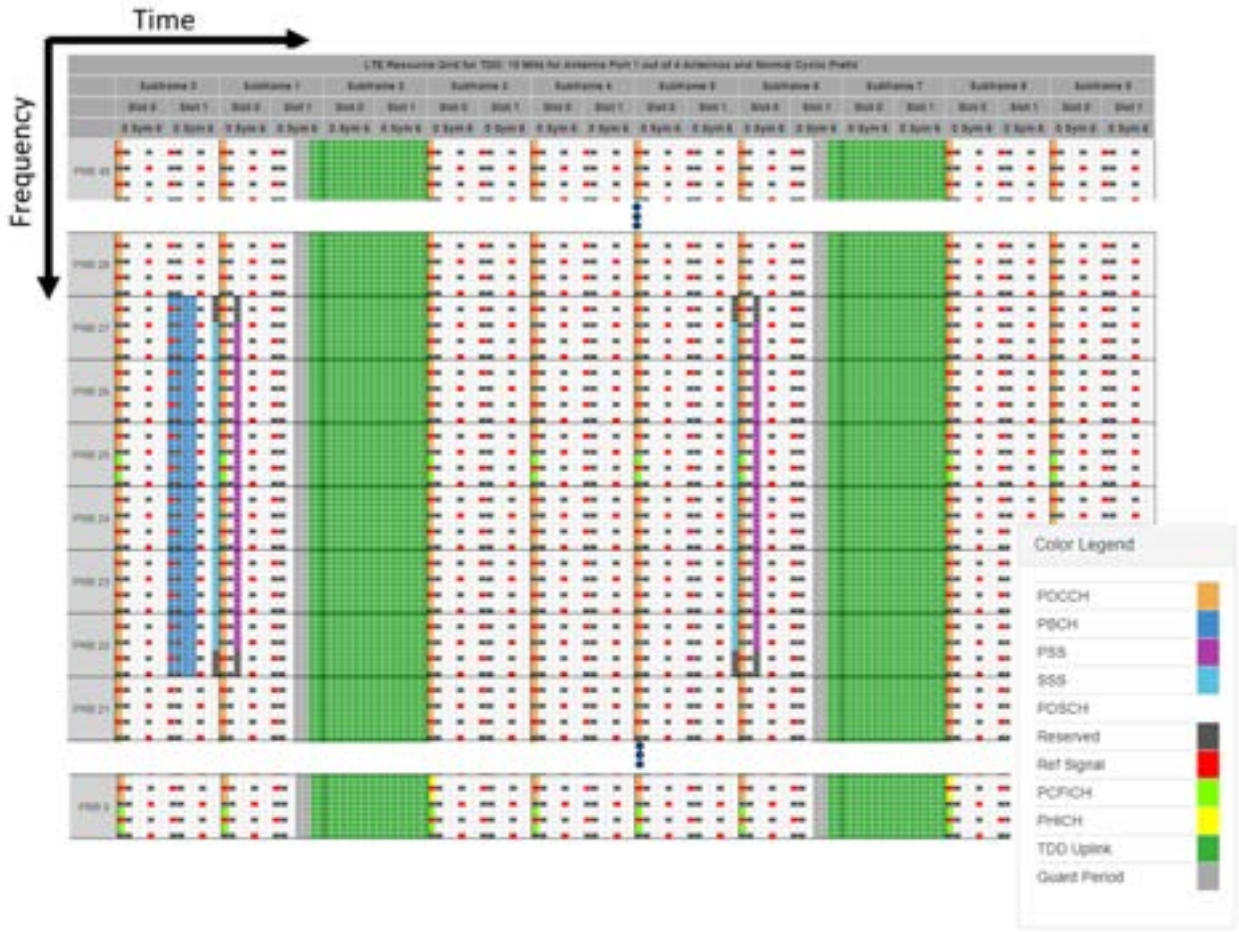

Figure 6.11: Resource grid for 3GPP TDD configuration 2, special subframe configuration 7 for a 10 Mhz channel (reference: dhagle.in/LTE)

The PFP traces can also be aligned across an entire day as seen in Figure 6.13, which shows this trend across longer periods of time. in Figure 6.13, the x-axis is "short-time" or the time duration of the PFP for the given capture while the y-axis is "long-time" or a waterfall of the captures throughout the day. The delta between the non-shared channel and purely shared channel symbol periods, as indicated by relative color changes, can give insights to how the channel is used throughout the day, possibly with respect to diurnal patterns.

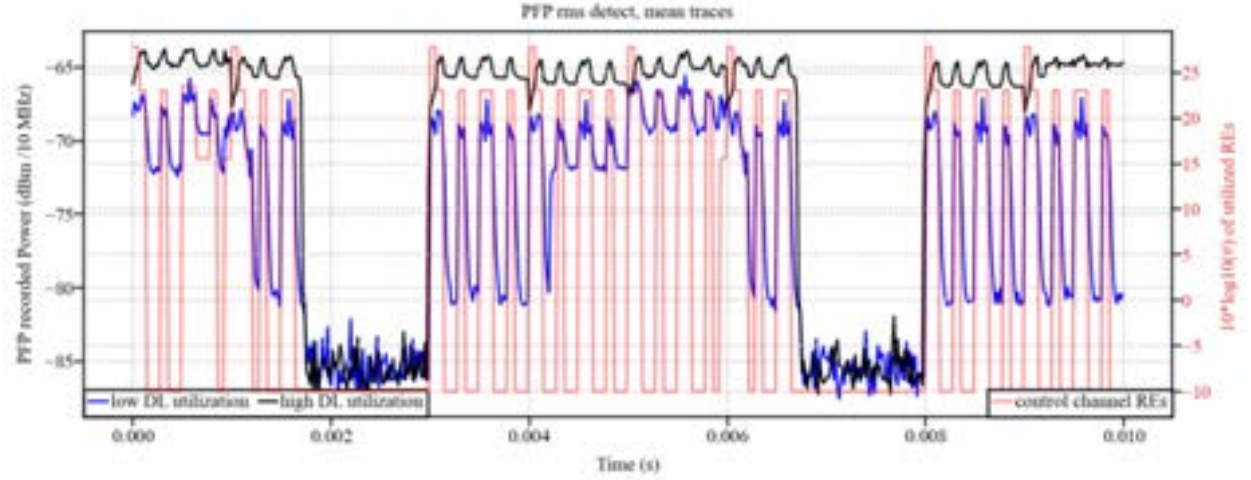

Figure 6.12: Aligning a symbol mapping of non-shared channel resources onto the PFP data product for a 10 MHz 3GPP allocation.

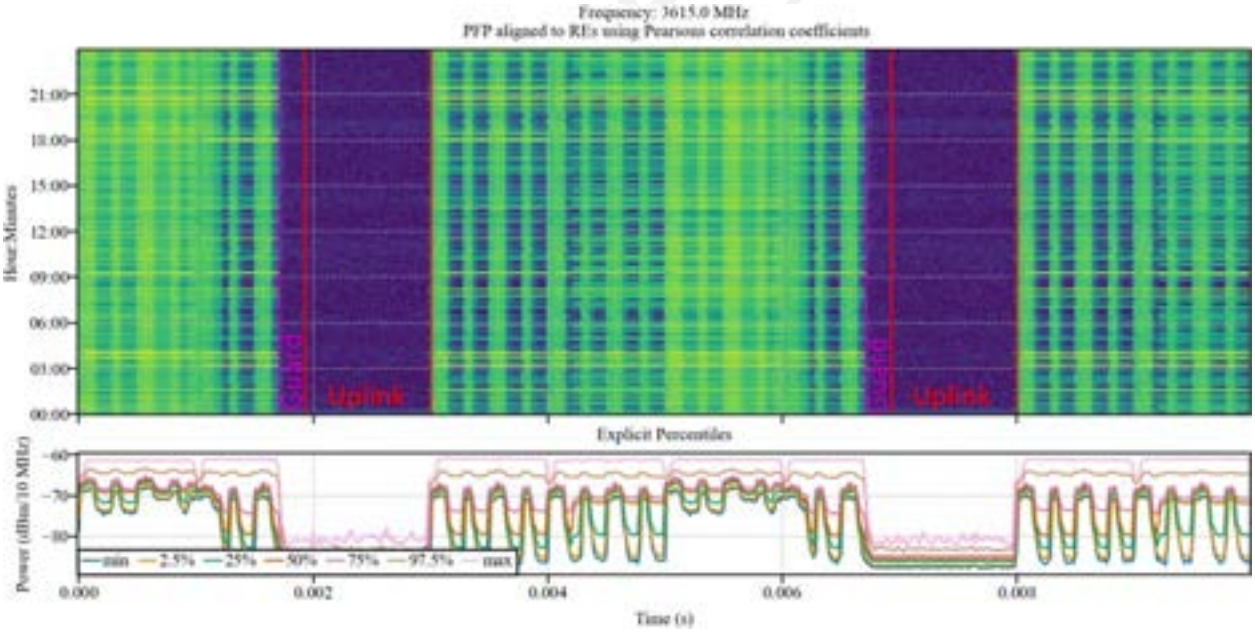

Figure 6.13: Aligned PFP traces as a waterfall through a 24 hour period.

Chapter 7

Conclusions

This test plan provides the framework and current approach to collect emissions in the Citizens Broadband Radio Service (CBRS) band in the vicinity of at least two Coastal Dynamic Protection Areas (DPAs) and provide data products to CBRS stakeholders to understand and assess the viability of CBRS spectrum sharing. Keep in mind that this plan is being developed in parallel with active site surveys and sensor deployments so there will be gaps and incomplete information in this test plan. National Advanced Spectrum and Communications Test Network (NASCTN) will detail any information not included in this test plan in the final publicly released test report. In the interim, any CBRS stakeholder can reach out to NASCTN for the most curernt information on sensor locations, data collection, and data products.

The NASCTN engineered and built custom sensors to ensure the required rigor, specifications, and tolerances were met during emission collection. The sensors were built to operate outdoors, over-time, with minimal hand-ons operations or maintenance. NASCTN leveraged existing National Telecommunications and Information Administration (NTIA) spectrum-characterization and occupancy sensing (SCOS) to collect and process the raw emission data into usable data products that were sent back to NASCTN for additional processing, storage, and dissemination to the CBRS stakeholders.

NASCTN used modeling and simulation to identify potential sensor locations as well as detailed analysis of selected locations to provide confidence the sensors would collect the needed and expected emissions. NASCTN deployed 2 sensors in the East Coast in the Norfolk, VA area with plans to deploy 1 more sensor in the East Coast and 4 sensors in the West Coat in the San Deigo, CA area. The intent of the sensor placement was to identify locations that capture emission close to Citizens Broadband Radio Service device (CBSD) deployments as well and sensors farther away from deployments to capture aggregate emissions.

The ultimate intent of this test plan is to provide a variety of data products that allow detailed analysis by incumbents and new operators to understand the growth and evolution of the CBRS ecosystem over time, the execution of the spectrum sharing paradigm, and identify anomalies that may warrant further investigation.

Appendix A

Sensor Architecture

A cursory description of the changes made to the sensor in the newest revision is given in Chapter 4. However, this appendix documents the entire sensor construction in the case that additional sensors are constructed. The overall sensor architecture remains similar to that of the previous Test Plan [2, Chapter 4] with some minor modifications to address issues that appeared once the prototype sensors were deployed.

A.1 Sensor Hardware Overview

The Sharing Ecosystem Assessment (SEA) sensor was designed to analyze radio frequency (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 design details on revision 1 of the production sensor.

A.1.1 RF System

Figure A.1 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 (designated 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 calibrations as described in Section A.1.3. Following the RF switch are an isolator (ISL) and a bandpass filter (BPF). The filter rejects undesired 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 variation 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.

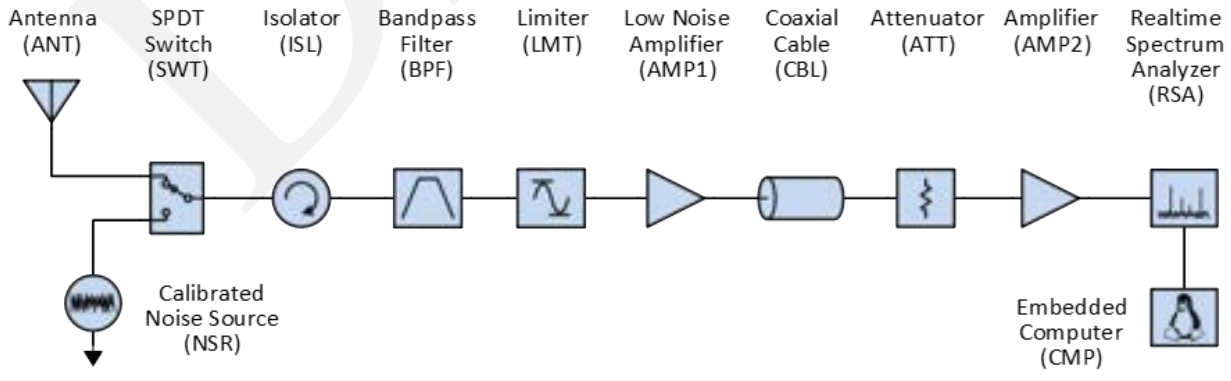

Figure A.1: Sensor RF System Block Diagram

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 and provides 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 is selected such that the sum of the cable insertion loss and attenuator insertion loss is a constant value, resulting in consistent RF performance regardless of the length of the cable used at a site 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. The RSA, while able to operate as a stand-alone and fully-featured test instrument, serves as the sensor's tunable RF digitizer. The RF signal at the input to the RSA 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 A.2 and A.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.

A.1.1.1 Antenna (ANT)

The sensor was designed to be used with an omnidirectional antenna to enable full 360° measurement in the azimuth plane. However, the sensor can easily accommodate directional antennas to support other measurement objectives. While this section details the omnidirectional antenna that is typically deployed, other antennas may be used with appropriate changes to the system calibration.

For typical applications the L-Comm HG3509U-PRO [10] antenna is used. It has 9 dBi gain and has a specified frequency range of 3400 to 3700 MHz. This antenna is omnidirectional in the azimuth plane and has a vertical 3 dB beamwidth of 11°.

Since CBRS systems are terrestrial systems, the antenna's directive elevation pattern is useful as neighboring CBSDs will be predominately located within its vertical 3 dB beamwidth.

The antenna is vertically polarized. Since CBSD devices typically leverage dual slant 45° polarized antennas, the measurement system may incur up to a 3 dB polarization mismatch loss. A more severe polarization mismatch loss would be incurred from horizontally polarized emitters, although horizontally polarized CBRS emitters are not typical.

A.1.1.2 Bandpass Filter (BPF)

A custom-designed cavity filter from Reactel, part number 8C7-3610-X180S11, 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. To minimize this impact, the filter was specified to have an insertion loss of less than 1 dB at the center frequency.

Data collected from previously deployed prototype sensors revealed that the sensors were vulnerable to strong C-band 5G emissions. C-band 5G begins at 3700 MHz, immediately adjacent to the CBRS band, and devices operating in this band are permitted to transmit at higher power levels than those operating in the CBRS band. The bandpass filter used in the prototype sensor was designed to have low insertion loss across the entire CBRS band, up to 3700 MHz, and offered little protection at the low end of the C-band. As a result, the prototype sensor sometimes became saturated by these strong C-band 5G signals while making measurements at the high end of the CBRS band. To combat this, the filter in the revision 1 production sensor was redesigned to provide greater suppression beginning at 3700 MHz. The

Table A.1: Measured Filter Response Characteristics

Center Frequency	3610 MHz
3 dB Bandwidth	160 MHz
Insertion Loss (3610 MHz)	-0.5 dB
CBRS Band Flatness (3550-3680 MHz)	0.4 dB
CBRS Band Upper Edge IL (3700 MHz)	-18 dB

Rejection	**Low**	**High**
>10 dB Rejection	<3524 MHz	>3695 MHz
>30 dB Rejection	<3508 MHz	>3712 MHz
>60 dB Rejection	<3458 MHz	>3752 MHz

trade-off to this solution is that the filter exhibits greater insertion loss at the upper end of the CBRS band, resulting in increased system noise figure and degraded measurement sensitivity at those frequencies.

The redesigned filter also extended the passband slightly lower in frequency, beyond the lower edge of the CBRS band, to enable operation at the upper end of the 3.45 GHz band (3450 MHz to 3550 MHz).

The new filters were measured using an Agilent N5230C vector network analyzer and the resulting frequency response for one unit is shown in Figure A.2. A summary of measured response characteristics is provided in Table A.1.

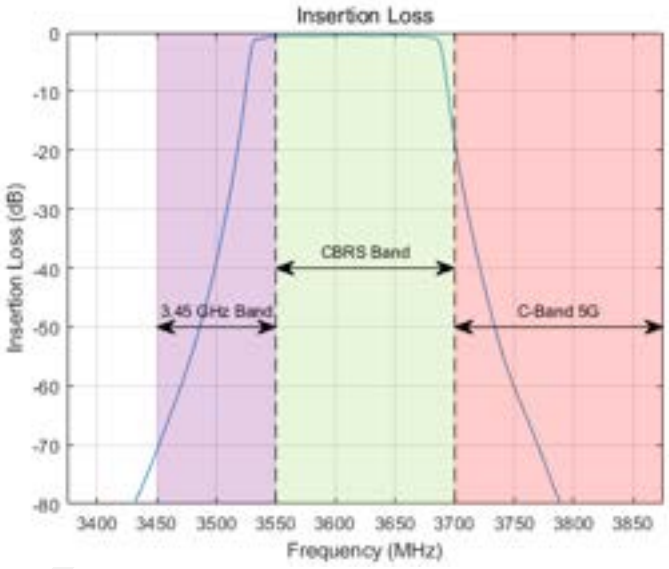

Figure A.2: Measured Filter Frequency Response

A.1.1.3 Low Noise Amplifier (AMP1)

The first LNA in the signal path is the LiConn LNA02004000A. This widebandLNA has specified operating frequency range of 0.2 GHz to 4.0 GHz. The typical gain at CBRS frequencies is 29 dB, as shown in Figure A.3, and a typical noise figure is 1.3 dB as shown in Figure A.4 [1]. The LNA is rated for a wide operating temperature range of -45° C to +85° C which simplifies the enclosure design by not requiring active heating or cooling.

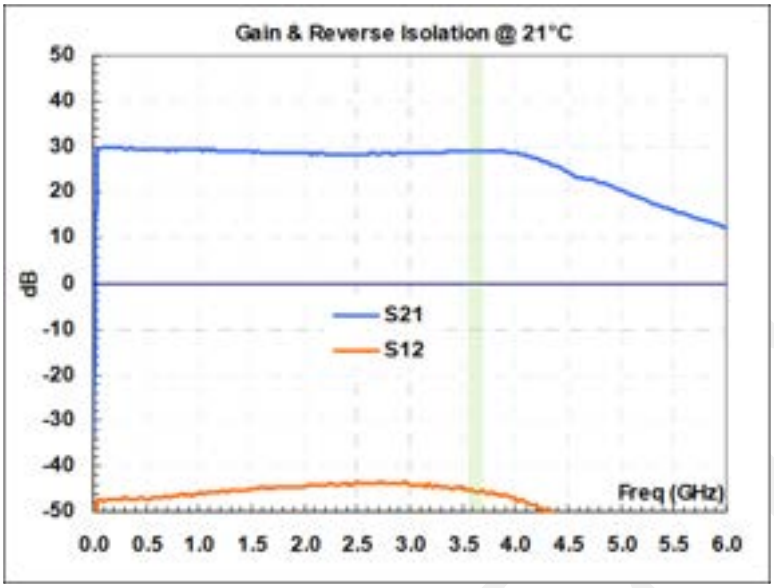

Figure A.3: LNA Gain and Reverse Isolation from Manufacturer Datasheet [1]

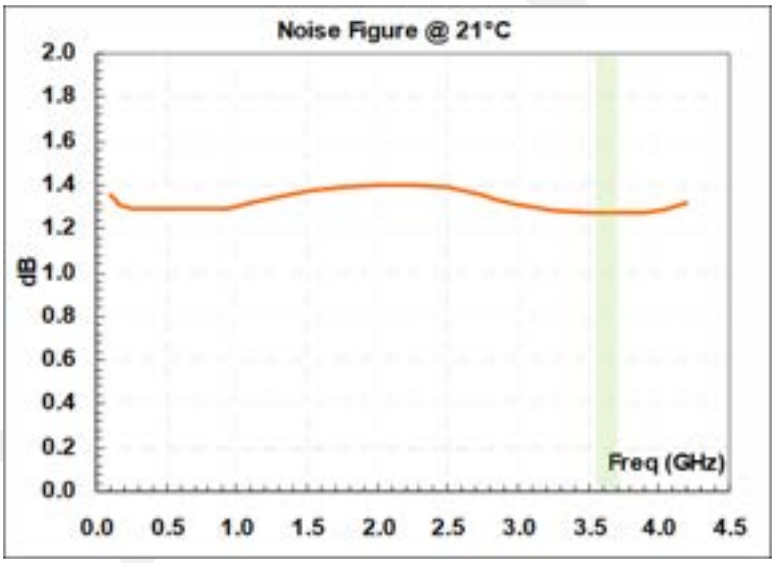

Figure A.4: LNA Noise Figure from Manufacturer Datasheet [1]

A.1.1.4 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. A number of candidates were identified and evaluated based on RFs performance, size, power consumption, operating temperature range and cost. Of those candidates, three were selected to undergo laboratory testing to examine RF performance, and ultimately the Tektronix RSA507A was selected.

The RSA507A realtime spectrum analyzer is tunable over the frequency range 9 kHz to 7.5 GHz and has an instanta-

neous 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 instantaneous 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 A.5 shows an example of the FFT analysis performed on the capture when the input frequency was 3605 MHz. The green marker indicates the input, and the red markers indicate 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 A.5 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.

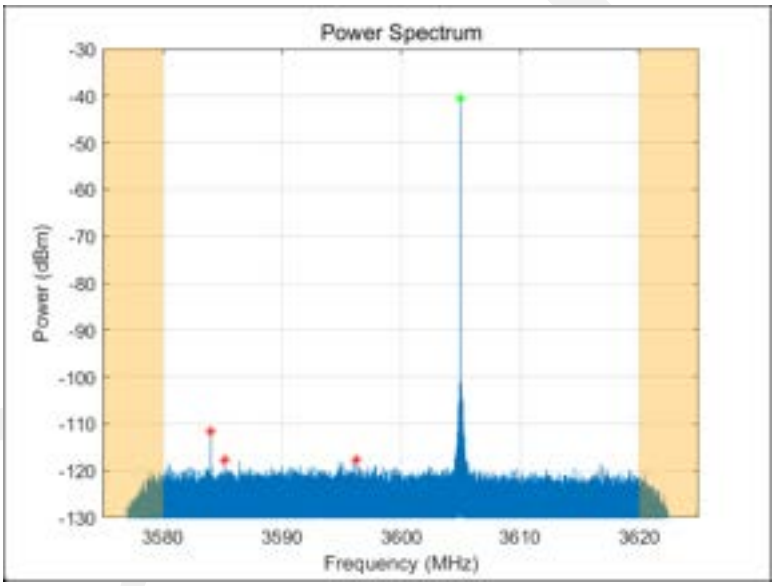

Figure A.5: RSA Power Spectrum with 3605 MHz CW Input Signal

Figure A.6 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.

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 with the internal preamplifier enabled, the RSA reference level was increased from -54 dBm to 0 dBm in 6 dB steps.

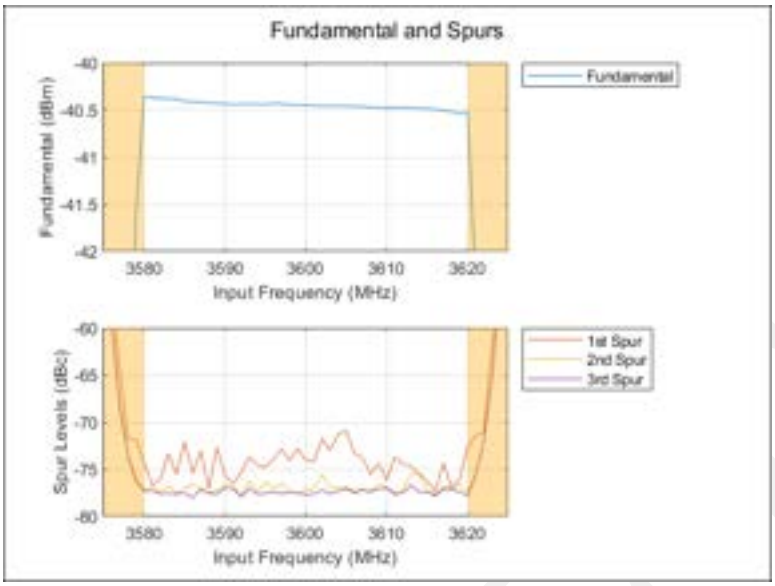

Figure A.6: RSA Fundamental Signal and Spur Levels

At all reference level settings, the RSAs 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 A.7. The minimum measured DANL of -155.6 dBm is within the range specified in the datasheet [11].

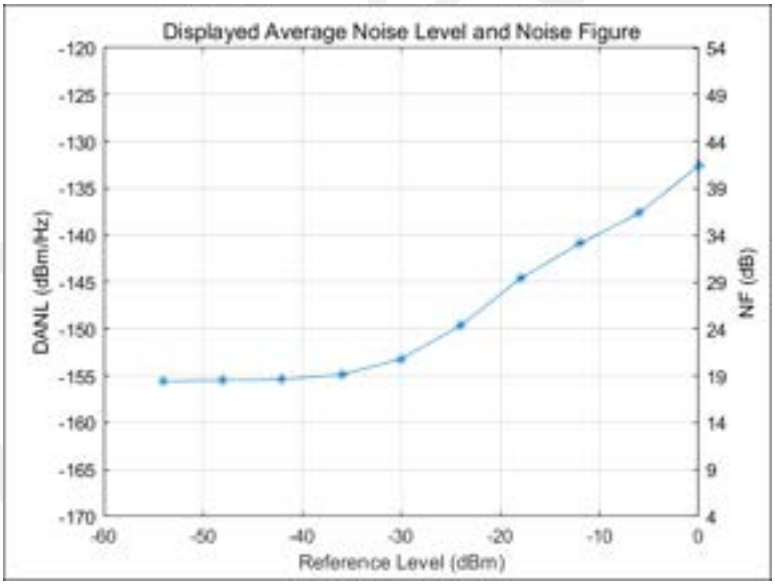

Figure A.7: RSA Displayed Average Noise Level and Noise Figure

1385 The RSAs equivalent noise figure (NF) is indicated on the right axis of Figure A.7. It is important to know the RSAs noise figure 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 A.8. 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 A.8 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 A.9.

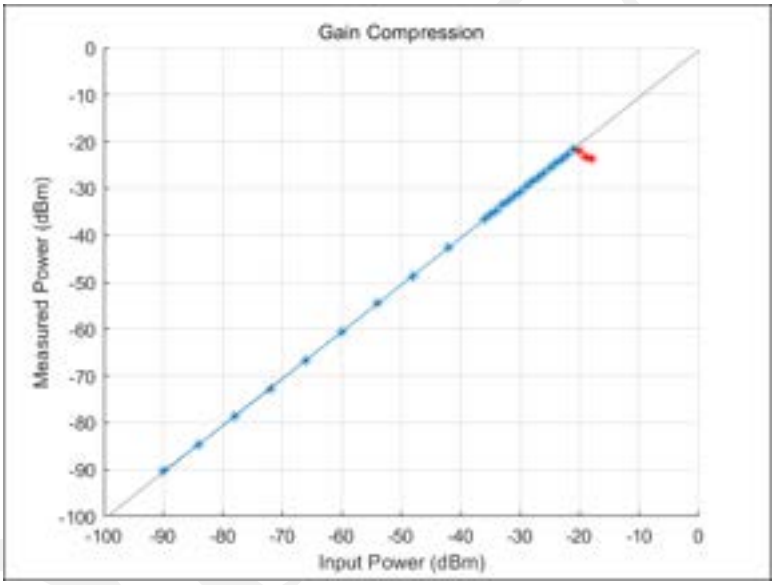

Figure A.8: RSA Gain Compression Measurement Example

A.1.1.5 Embedded Computer (CMP)

Processing of the IQ data is performed on a SimplyNUC NUC12TZi7 small form factor computer. This computer has an 12th 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 5.2.

A.1.2 Cascade RF Analysis

The system cascade analysis has been generated to characterize the RF signal path from the input of the preselector to the signal processor. The goals of this analysis is to characterize the signal levels in the signal processing chain and to provide a low noise figure by adding a minimal amount of noise to the received signal. The receive chain gain is designed to maintain noise figure of the receive signal while also allowing for a high dynamic range of the measurement. The components of the receiver chain need to be selected to address overall system constraints and

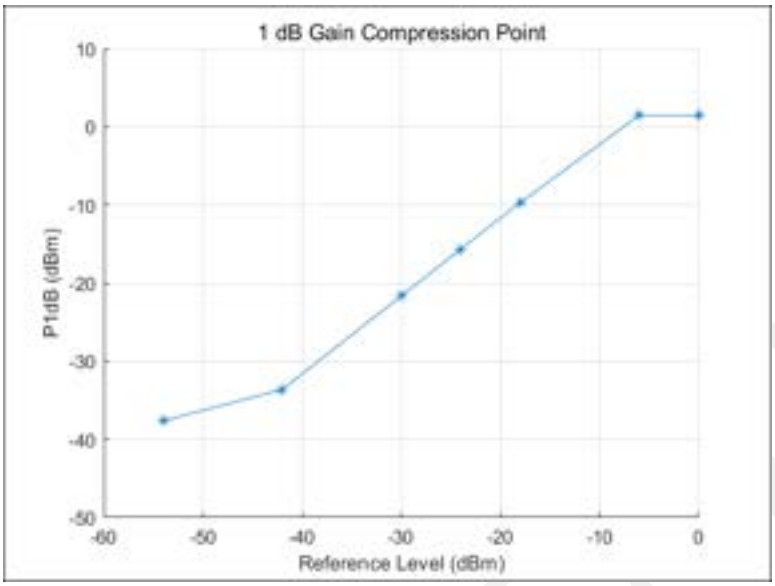

Figure A.9: RSA 1 dB Gain Compression Point

goals including:

- 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, device linearity, maintaining a low noise figure, and ensuring the signal level is within the dynamic range of the signal analyzer. The first step in the design was to perform a cascaded system analysis of the components used in the receive chain based on information provided in the component data sheets.

In the system cascade, the noise figure is dominated by loss of the components that are in front of the first LNA gain stage. The design of the preselector includes a calibration switch and a bandpass 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 includes a second LNA to provide additional signal gain in front of the spectrum analyzer that helps maintain noise figure through the cabling between the preselector and the signal processing unit. This allows for a long cable to be placed between the preselector and the signal processing unit for flexibility in installation locations. An attenuator is included in the signal processing unit to ensure flexibility in the event that a longer cable is required between the preselector and signal processing unit. A full system block diagram is shown in Fig. A.10. The gain of full system is 32.1 dB and the noise figure is 4.9 dB as shown in Fig. A.11. The dominant factor in this overall noise figure is due to loss from the components before the first LNA stage.

Finally, utilizing the cascaded gain and noise figure for the system is shown versus frequency in Fig. A.12 demonstrating performance across the entire CBRS band. The front-end bandpass filter is designed to provide attenuation at the upper CBRS band edge of 3700 MHz. The bandpass filter helps to attenuate C-Band Fifth Generation of cellular mobile communications (5G) emissions in the 3700 to 3980 MHz band. The filter mainly impacts the noise figure of

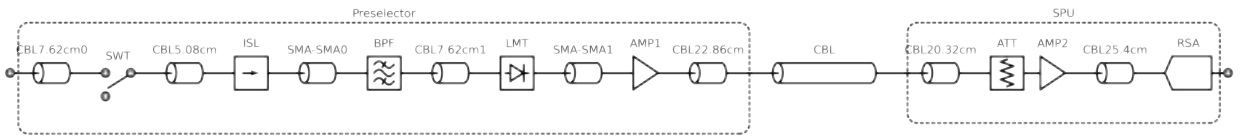

Figure A.10: Full system cascade block diagram illustrating the components and interconnects that are considered in estimating the cascaded gain and noise figure.

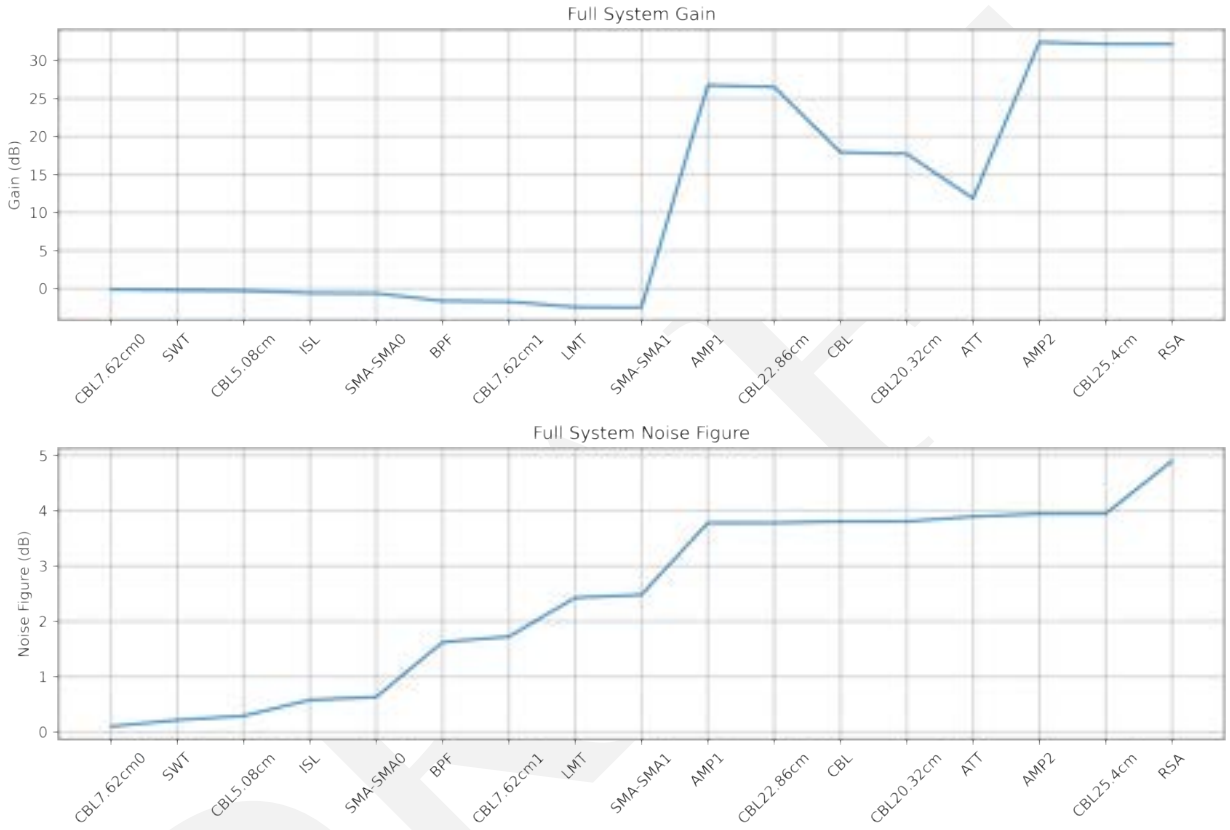

Figure A.11: Cascaded gain (top) and noise figure (bottom) of the full system showing the contribution of each component to the overall system gain of 32.1 dB and noise figure of 4.9 dB.

the upper two channels of the CBRS band due to the additional insertion loss while protecting the signal analyzer from being overloaded.

A.1.3 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 of this effect is shown in the NTIA TR-20-548 [12]. In this report, the gain of the measurement sensor shown in Figure A.13, 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

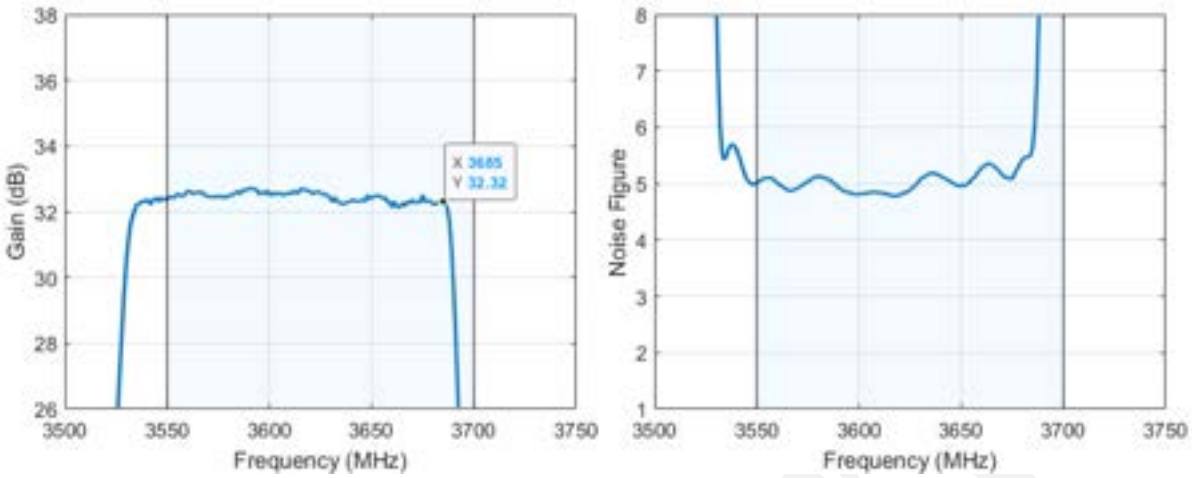

Figure A.12: Cascaded gain and noise fgure prediction of the system over the entire CBRS band

for by leveraging an on-board calibration source.

After bench and initial field calibrations described in Section A.4, this calibration with the on-board noise diode will occur periodically. Initially, it will occur at 60 minute intervals. However, as calibration actions detract from measurement time, the calibration variability will be analyzed and the frequency of on-board calibrations may be reduced in the future.

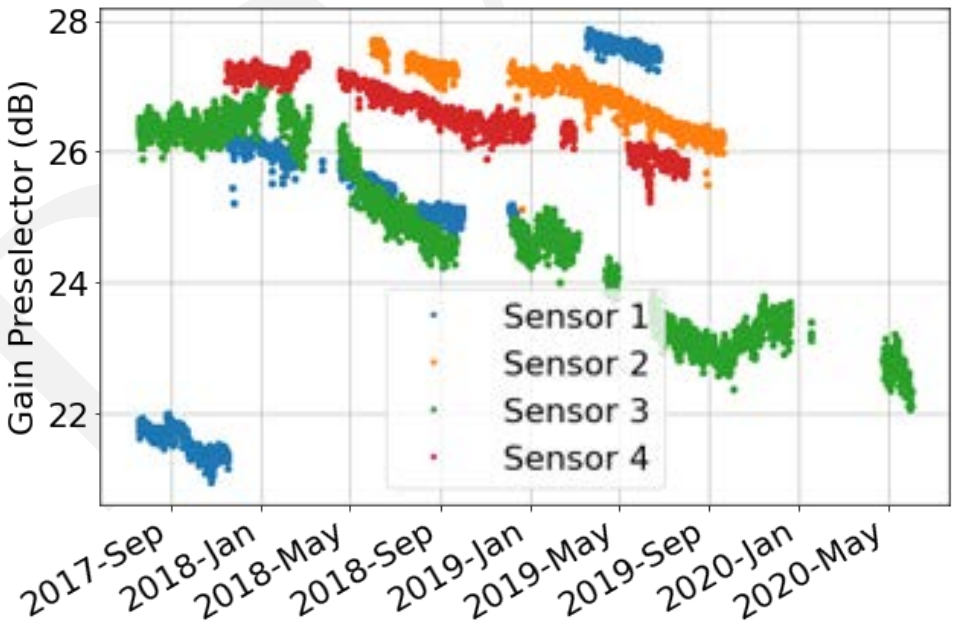

Figure A.13: Gain drift shown of measurement sensor leveraged in NTIA's TR-20-548 Technical Report.

A.2 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 first gain stage (designator AMP1 in Figure A.1), 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, minimizing antenna cable loss. 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. Like the preselector, the SPU is assembled into a weatherproof enclosure to enable deployment in areas exposed to the weather. With the power dissipation of the internal components and lack of air circulation, a thermo electric heater/cooler (TEC) was included to prevent internal temperatures from exceeding operating limits. Additionally, the TEC has heating capability to prevent the internal temperature from falling below the operating range. The SPU may be installed up to 30 m away from the preselector, and for installation in climate-controlled areas, the TEC may be disabled or removed. Further details on preselector and SPU thermal requirements and performance are provided in Sections A.2.1.1 and A.2.2.1, respectively.

Figure A.14 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.

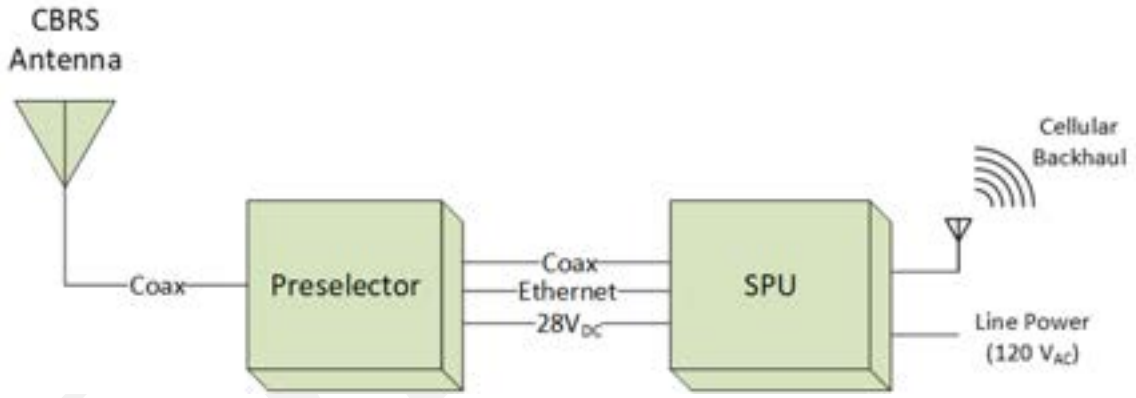

Figure A.14: SEA Sensor Implementation Diagram

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

A.2.1 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 A.15 and component part numbers are provided in Table A.2. Note that Table A.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 A.15) allows the RF switch (SWT) state, power to the LNA (AMP1), and power to the noise source (NSR) to be controlled remotely. 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.

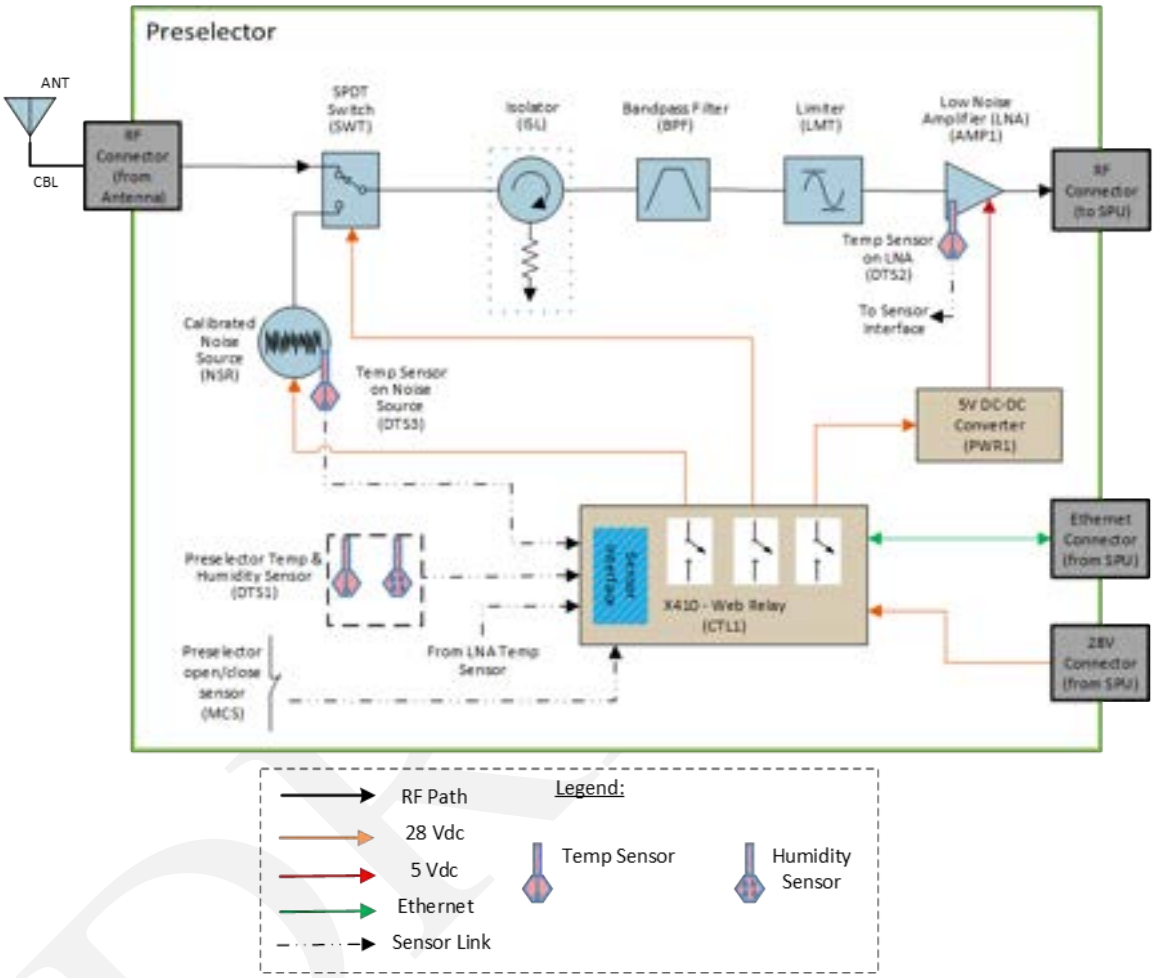

Figure A.15: Preselector Diagram

The preselector is assembled in a 35.6 cm x 30.5 cm x 17.8 cm fiberglass-reinforced polycarbonite 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 with components annotated is shown in Figure A.16. The bottom side of the enclosure is modified with a connector panel, shown in Figure A.17, 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.

Table A.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, 3530-3690 MHz	Reactel	8C7-3610-X180S11	-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
MCS1	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

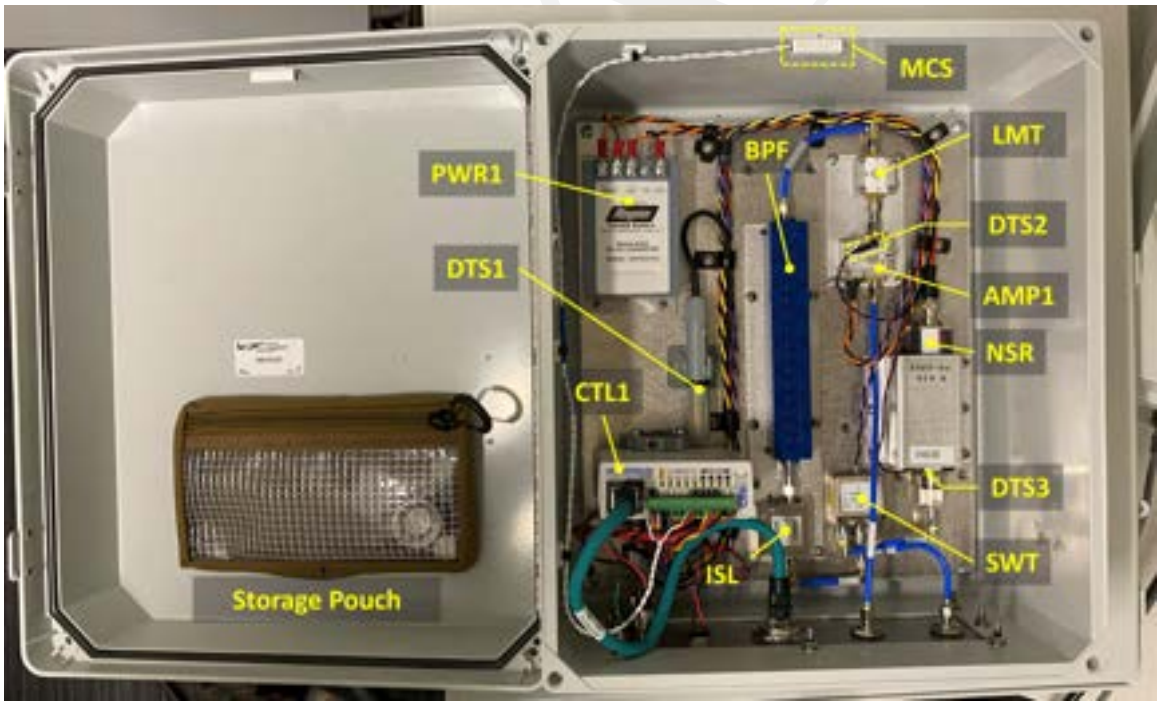

Figure A.16: Preselector Assembly

Figure A.17: Preselector Connector Panel

A.2.1.1 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 A.18 shows the 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.

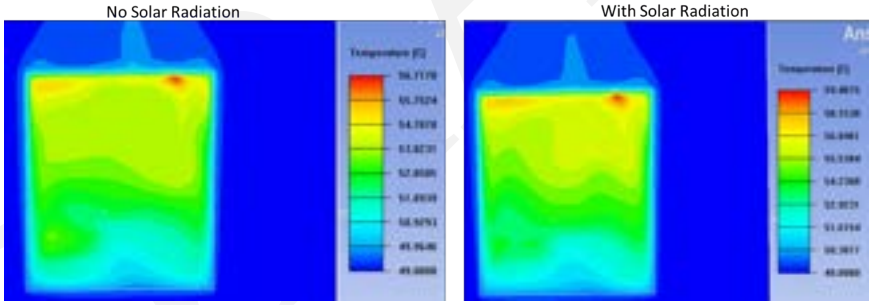

Figure A.18: Thermal Analysis Results Without Solar Radiation and With Solar Radiation

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 A.16

- 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 A.19 shows the recorded temperatures.

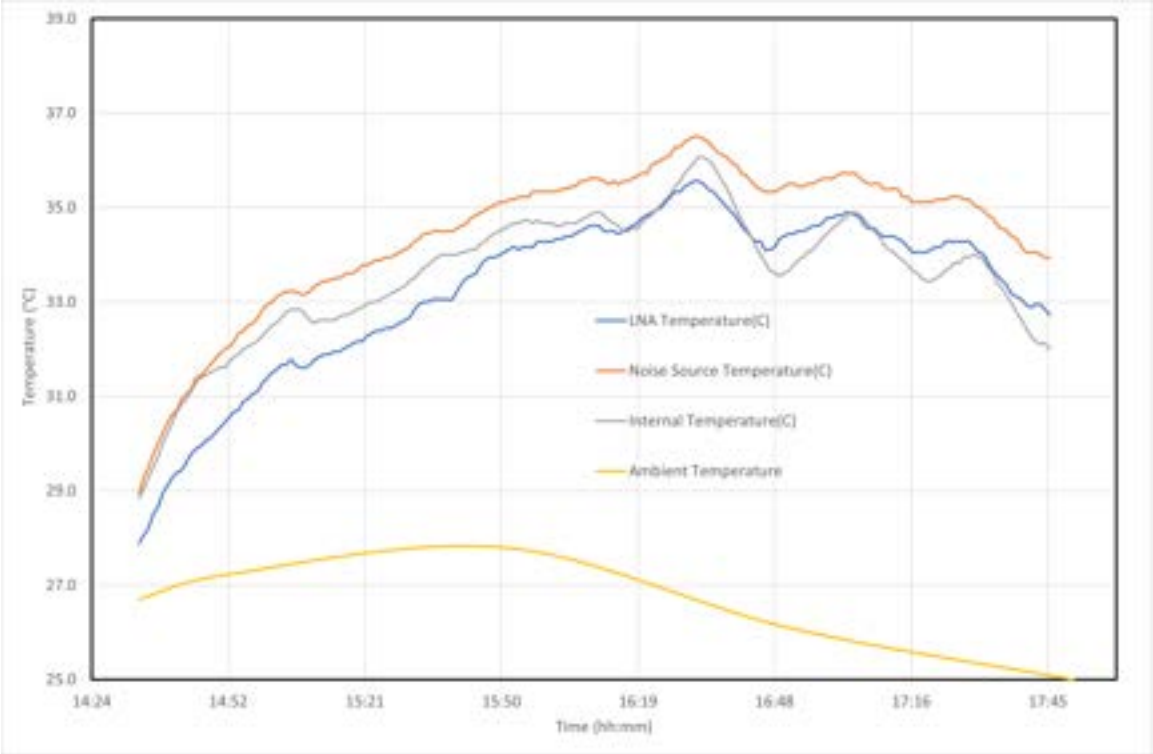

Figure A.19: Preselector Outdoor Thermal Test Results

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.

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 A.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.

Temperature data was collected for fielded sensors from 7 August to 18 August 2023. Ambient temperature was obtained from historical data at the closest weather stations. The maximum ambient temperatures recorded during this time period was 34° C near Hampton University at Newport News Williamsburg International Airport Station and 36° C near Norfolk International Terminal (NIT) at Norfolk International Airport Station. The maximum temperatures recorded inside the preselector enclosure were 44.4° C at Hampton University and 46.7° C at NIT. Figures A.20 and A.21 show the internal preselector temperatures and weather station temperatures.

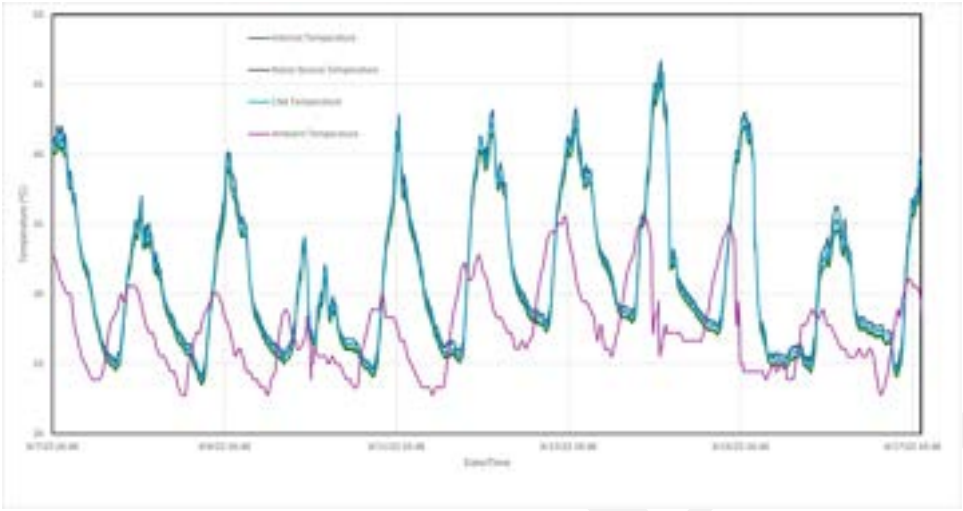

Figure A.20: Temperatures at NIT in August 2023

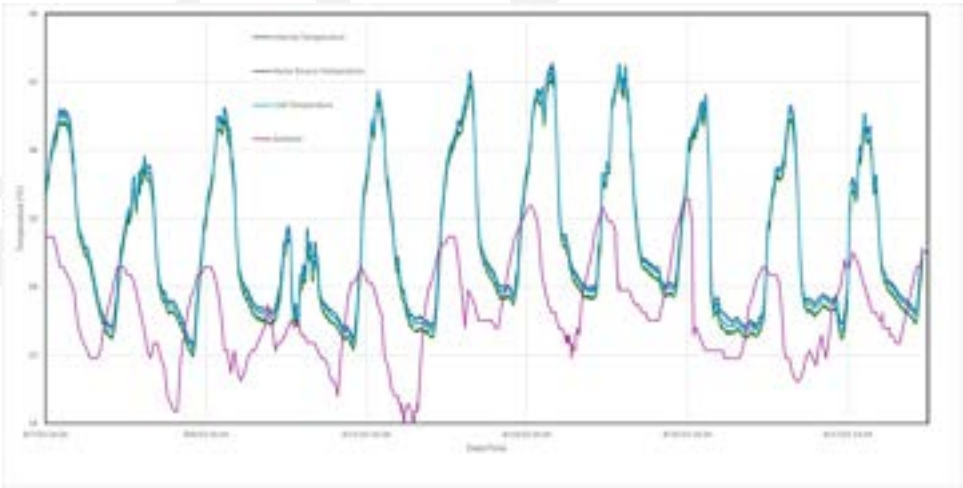

Figure A.21: Temperatures at Hampton University (HU) in August 2023

A.2.2 SPU Assembly

Figure A.22 shows an overall block diagram of the SPU. A full listing of component part numbers referenced in the following paragraphs is provided in Table A.3.

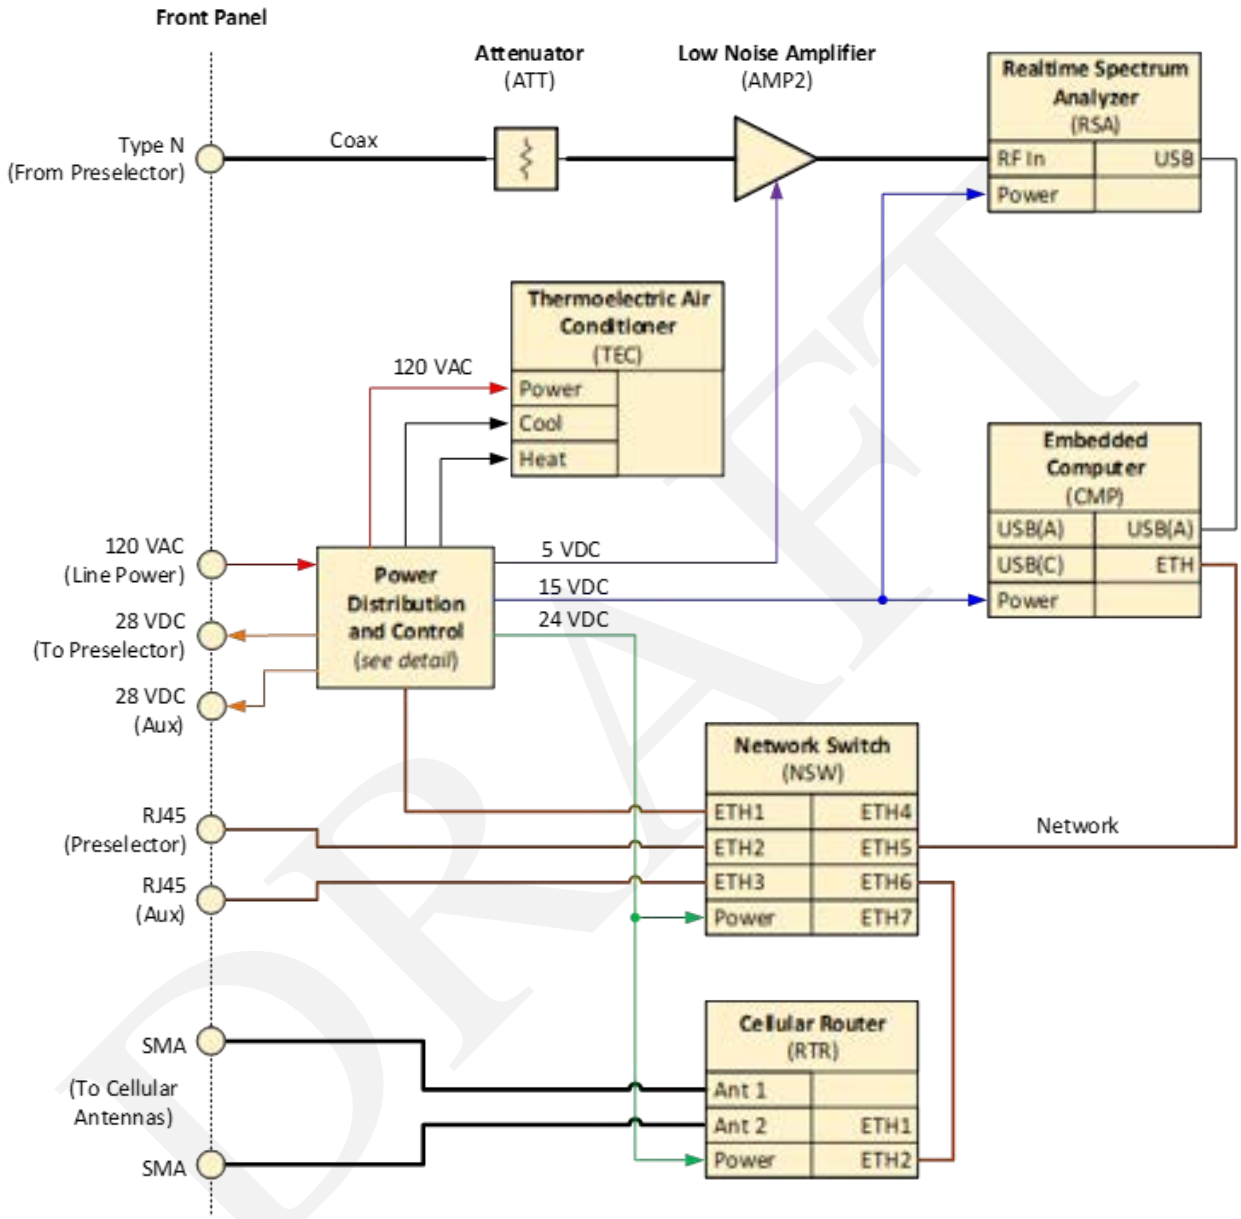

Figure A.22: SPU Overall Block Diagram

The SPU assembly contains the remaining RF components which are not integrated into the preselector, namely the
1530 attenuator (ATT), second amplifier (AMP2), and RSA. Additionally it contains the embedded computer (CMP) where data collection and processing are performed and various power management and system control components. Further details on the power distribution and control hardware are provided in Figures A.23 and A.24, respectively. Finally, the SPU assembly contains a TEC and heater, used to ensure that internal temperatures do not exceed component operating temperature limits. Further details regarding temperature control can be found in Section A.2.2.1.

Finally, the SPU contains a cellular router (RTR) to enable remote access to the sensor data products as well as remote monitoring and administration of the sensor. This mobile networking strategy enables sensor deployments in locations where wired network connectivity is not available but cellular coverage exists. A Digi IX20 LTE Router was chosen, as it provides the required connectivity and also serves as a network router which, along with the network switch (NSW), enables communication among the various network-enabled components in the sensor. Finally, these cellular routers are capable of being linked to a private network.

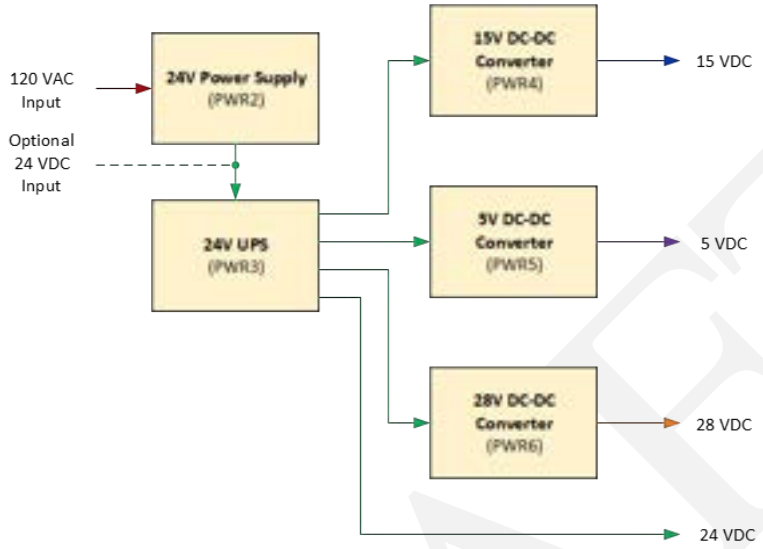

Figure A.23: SPU Power Distribution Block Diagram

The power management components in the SPU generate the required supply voltages and distribute power to the entire sensor, including the preselector. As shown in Figure A.23, the sensor is powered from a standard 120 VAC source. Inside the SPU the supply voltage is immediately converted to 24 VDC by a power supply (PWR2), which is then backed up by a DC uninterruptible power supply (UPS) system (PWR3). The 24 VDC output of the UPS supplies power to the rest of the sensor, except for the TEC.

With a simple reconfiguration, the sensor can be powered by an external 24 VDC source. In this configuration, the 24 VDC power supply (PWR2) is disabled and the UPS (PWR3) receives power from external 24 VDC source. It is important to note, however, that 120 VAC is required to operate the thermoelectric air conditioner (TEC). Therefore the sensor can only be powered by 24 VDC in environments where temperatures will be moderate.

The 24 VDC output of the UPS directly powers several components including the network switch (NSW) and cellular router (RTR). The 24 VDC output is also converted to several other DC voltages. A 15 V DC-DC converter (PWR4) generates 15 VDC for the embedded computer (CMP) and RSA. A 5 V DC-DC converter (PWR5) generates 5 VDC for the LNA (AMP2). Finally, a 28 V DC-DC converter (PWR6) generates the 28 VDC supply voltage needed for the preselector.

A network-enabled controller (CTL2) along with several expansion modules (CTL3 to CTL5) perform a variety of system control and monitoring functions. Details of these controllers are shown in Figure A.24. The main controller

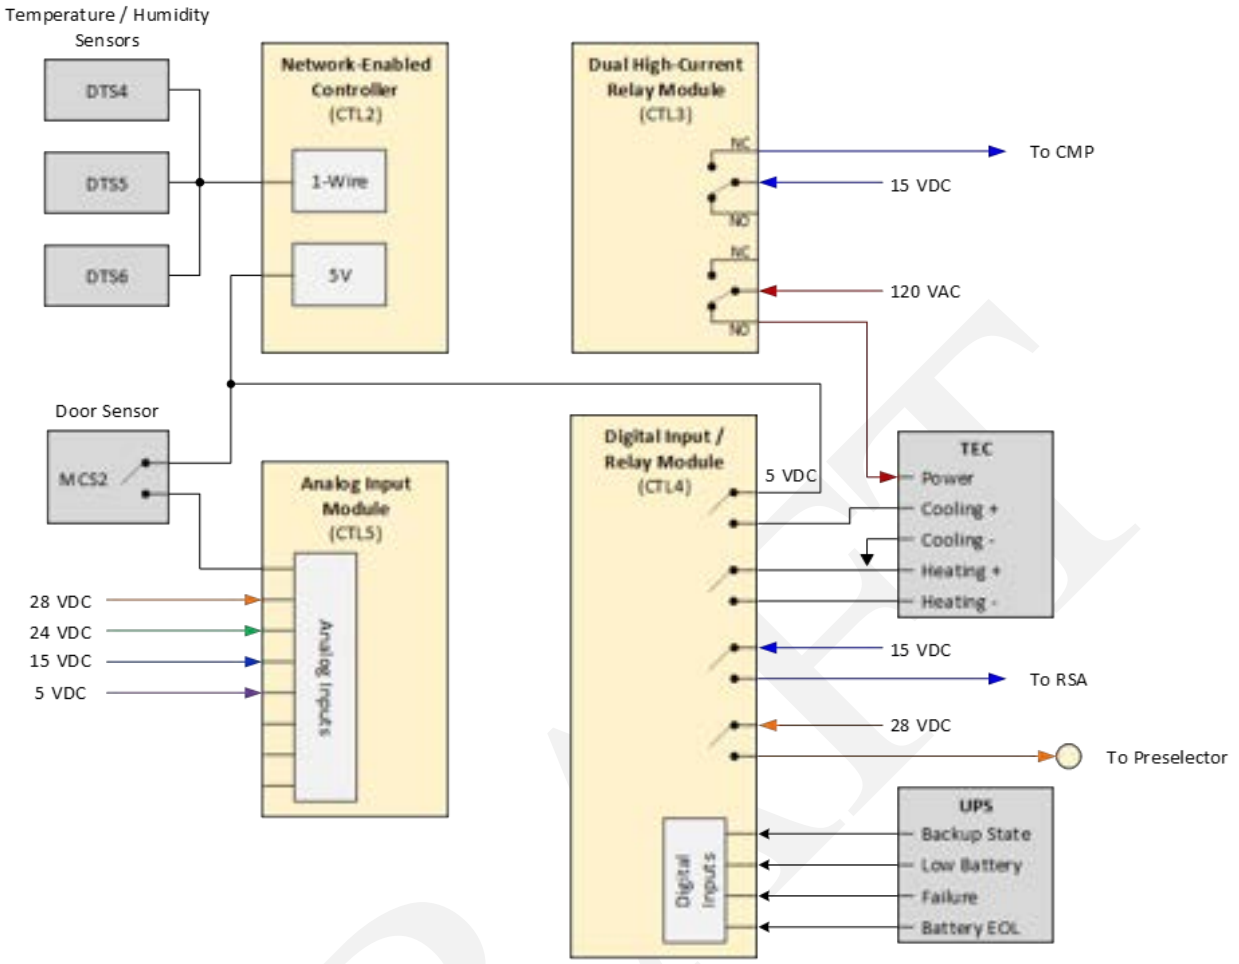

Figure A.24: SPU Control and Monitoring Block Diagram

(CTL2) enables remote monitoring of several temperature and sensors, and also interfaces with the expansion modules. These sensors include a temperature sensor placed near the air intake of the TEC, a temperature sensor placed near the exhaust, and a combined temperature/humidity sensor placed in a central area of the assembly.

The dual high-current relay module (CTL3) allows the DC voltage to the computer (CMP) to be enabled or disabled. While power should always be supplied to CMP during normal operation, it's important to have the ability top remotely toggle power in the event of a system failure. The high-current relay module also allows the 120 VAC supply voltage to the TEC to be enabled or disabled. This allows the TEC to be turned off during periods of moderate temperature.

The digital input / relay module (CTL4) serves a number of purposes. Relays in this module are used to control the supply voltages to the RSA and preselector. As with CMP, during normal operation power to both the RSA and preselector will always be enabled, but the ability to remotely toggle power is important for troubleshooting. These relays are also used to enable the cooling and heating modes of the TEC. The digital inputs on CTL4 are used to monitor several UPS status and health indicators. These indicators alert the system when a power failure has occurred, when the backup battery is at a low state of charge, or when a UPS failure has been detected.

Finally, the analog input module (CTL5) is used to remotely monitor the various DC voltages in the system. It also

Table A.3: SPU Components

Designator	Description	Manufacturer	Part Number	Temp Range
AMP2	Low noise amplifier, 500 MHz - 8 GHz	Mini-Circuits	ZX60-83LN-S+	-40 to +85 °C
ATT	Attenuator, 6 dB, DC - 18 GHz	Mini-Circuits	BW-S6W2+	-55 to +100 °C
CMP	Small form factor computer, 12th Gen Intel® Core™ i7 processor, 32 GB RAM	SimplyNUC	NUC12TZi7	0 to +60 °C
CTL2	Network-enabled I/O controller	ControlByWeb	X400-I	-40 to +65.5 °C
CTL3	Dual high-current relay expansion module	ControlByWeb	X-11s	-40 to +65.5 °C
CTL4	Four relay + four digital input expansion module	ControlByWeb	X-17s	-40 to +65.5 °C
CTL5	Eight analog input expansion module	ControlByWeb	X-22s	-40 to +50 °C
DTS4	Digital temperature and humidity sensor	ControlByWeb	X-DTHS-P	-40 to +80 °C
DTS5 DTS6	Digital temperature sensor	ControlByWeb	X-DTS-U	-55 to +125 °C
MCS2	Magnetic contact closure sensor	ControlByWeb	MPS	-40 to +100 °C
NSW	Ethernet switch, Gigabit, 8-port	TRENDnet	TI-G80	-40 to +75 °C
PWR2	Power supply, 100 to 240 VAC input, 24 VDC output, 240 W	Omron	S8VK-X24024A-EIP	-40 to +70 °C
PWR3	Uninterruptible power supply, 24 VDC, 240 W, 3.2 Ah Li-ion battery	Omron	S8BA-24D24D240LF	0 to +55 °C
PWR4	DC-DC converter, 24 VDC input, 15 VDC output, 10 A	Traco	TEP 150-2413WI	-40 to +75 °C
PWR5	DC-DC converter, 24 VDC input, 5 VDC output, 2.5 A	Acopian	24EB5E250	-20 to +71 °C
PWR6	DC-DC converter, 24 VDC input, 28 VDC output, 0.5 A	Acopian	24EB5E250	-20 to +71 °C
RSA	Portable USB real-time spectrum analyzer, 9 kHz to 18 GHz	Textronix	RSA507A	-10 to +55 °C
RTR	4G LTE Router	Digi	IX20	-40 to +70 °C
TEC	Thermoelectric 400 BTU and heater, external controller option	EIC	AAC-140C-4XT-HC-RC	-20 to +60 °C

monitors a magnetic contact switch that indicates when the door of the SPU enclosure has been opened.

The SPU is assembled in a 63 cm x 53 cm x 25.5 cm fiberglass-reinforced polycarbonate weatherproof enclosure with a NEMA 4X / IP67 rating. As with the preselector, this enclosure was selected to allow deployment in extreme environmental conditions including temperature, humidity, and solar radiation. Due to the heat generated by the internal components and the lack of air flow imposed by the sealed enclosure, a TEC is included to prevent internal temperatures from exceeding operating limits. The TEC has heating capability as well to prevent the internal temperature from falling below the operating range during extreme cold temperatures. Further details on thermal control and testing are provided in Section A.2.2.1.

All RF and electrical components are mounted on series of aluminum brackets and plates, which are then mounted to a main base plate. A photo of the internal assembly is shown in Figure A.25. The bottom side of the enclosure is modified with a connector panel, shown in Figure A.26, with feed-through connectors for the RF input, cellular antennas, Ethernet and 28 VDC power output to the preselector. A power cord attached to the panel connects to a 120 VAC source, and a DC input connector can be used to optionally supply power to the SPU from an external 24 VDC source.

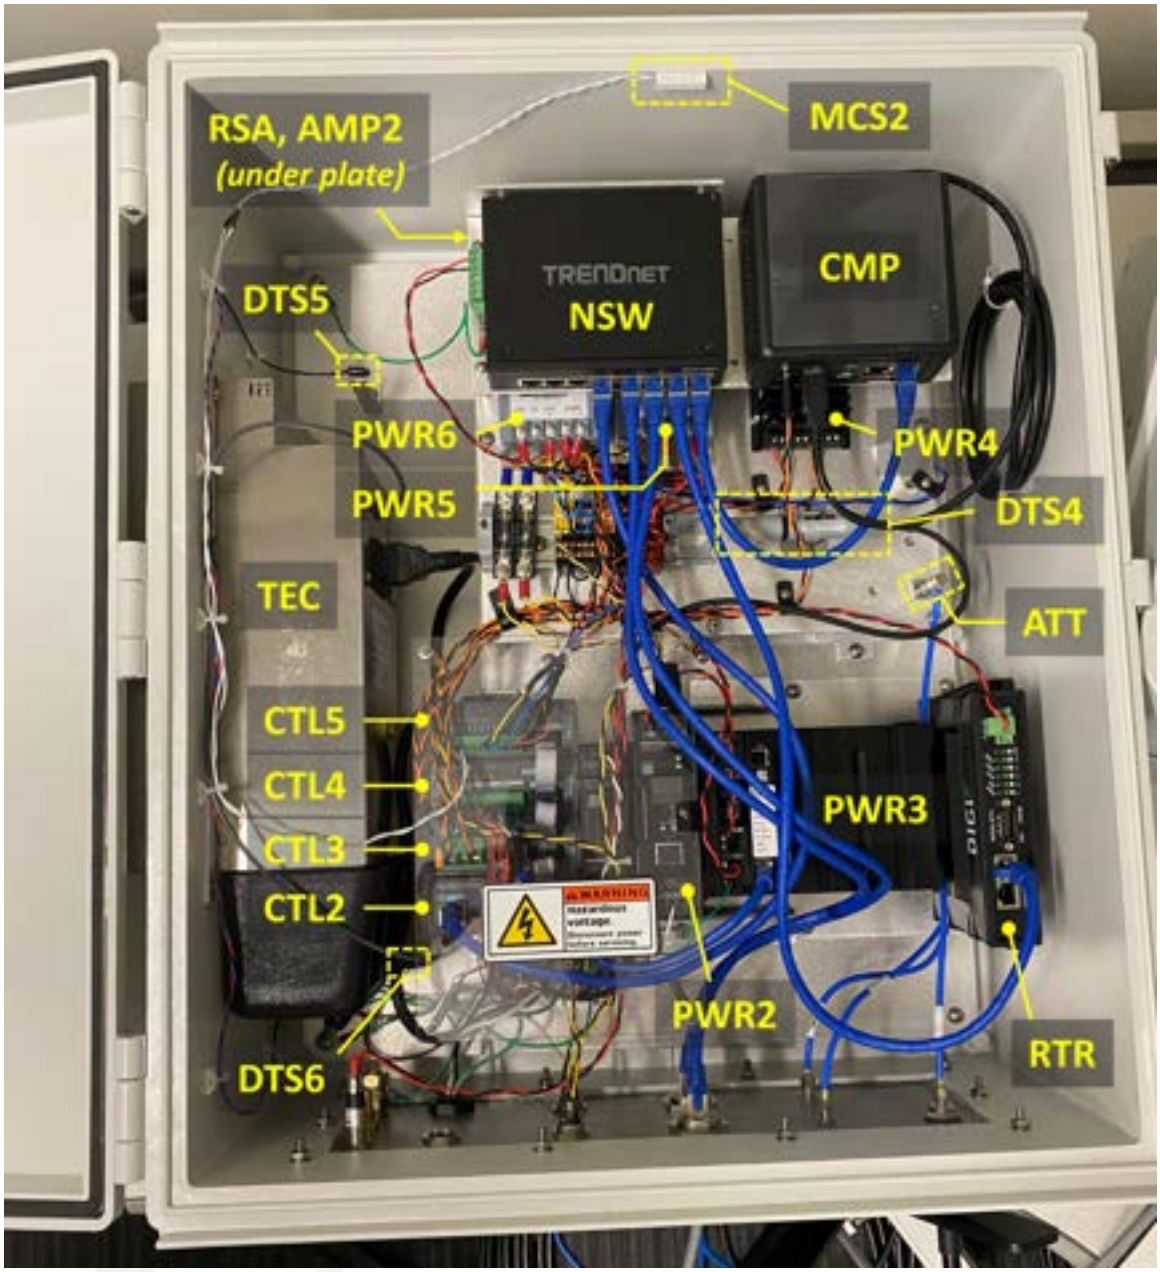

Figure A.25: SPU Assembly

A.2.2.1 SPU Thermal Control and Testing

A 400 BTU TEC with heater is installed in the SPU enclosure to control internal air temperature. The TEC cools the enclosure without the need for vents, therefore keeping the enclosure sealed and free from dust and water intrusion. When the TEC power is on, but cooling or heating is not enabled, the internal fans turn on and circulate air within the enclosure. When the cooling capability is enabled, the external fans turn on in order to remove heat from the TEC. The TEC is controlled by the controller (designator CTL2) and several associated expansion modules (CTL3 and CTL4). The controller monitors the temperature inside the enclosure with sensor DTS4, and enables TEC power, cooling mode, and heating mode when the temperature reaches certain setpoints. The SPU was tested in a thermal

Figure A.26: SPU Connector Panel

Table A.4: TEC Control Temperature Setpoints

	Cold Temperature Setpoints		Warm Temperature Setpoints	
	TEC Power	**Heating Mode**	**TEC Power**	**Cooling Mode**
Enable	<8 °C	<5 °C	>42 °C	>45 °C
Disable	>10 °C	>7 °C	<40 °C	<43 °C

chamber at MITRE Bedford to validate the design and determine the TEC controller setpoints. Table A.4 lists the setpoints that were used during this test.

Figures A.27 through A.29 show temperature plots from three thermal tests. Temperatures in three locations inside the SPU and the air temperature inside the thermal chamber are plotted as a function of time. The locations are marked as reference designators DTS4, DTS5, and DTS6, and their locations are shown in Figure A.25. Figure A.27 shows the temperatures with the TEC powered off at ambient temperatures of 0° C, 10° C, and 20° C. The chamber temperature was then increased to 30° C, but the test was terminated before the internal SPU air temperature reached 45° C to avoid approaching the maximum operating temperature of some of the components. This test showed that the SPU can operate in ambient temperatures from 0° C to 20° C without the TEC. During operation, the internal air temperatures stabilize at approximately 10° C to 18° C above ambient temperature. However, those numbers will increase when the enclosure is exposed to solar radiation.

Figure A.28 shows a high temperature test. The chamber temperature was set to 43° C. At 210 minutes into the test, the chamber temperature was ramped up to 45° C. At an ambient of 43° C, the cooling mode cycled on and off since the cooling turns on at an air temperature of 42° C and off at 40° C. The internal air temperature hit a peak of 51° C when the cooling was off, but quickly decreased as soon as the cooling turned on. Once the chamber temperature was increased to 45° C, the cooling remained on continuously, and the maximum internal SPU temperature remained constant at approximately 44° C. The thermal tests do not include solar radiation. As the internal temperature will increase when the enclosure is exposed to solar radiation, a solar shield design will be investigated and implemented where needed.

Figure A.29 shows a low temperature test. The chamber was set to -5° C for 98 minutes, then ramped to -15° C and held for 150 minutes, then set to -20° C and held until the end of the test. After 45 minutes at an ambient temperature of -15° C, the TEC power turned on. Five minutes later, the heater switched on and then switched off and back on until the chamber temperature was set to -20° C, at which point, the heater remained on constantly. Based on additional temperature data collected from the first fielded sensor at at Green Mountain Mesa, the team decided to adjust the

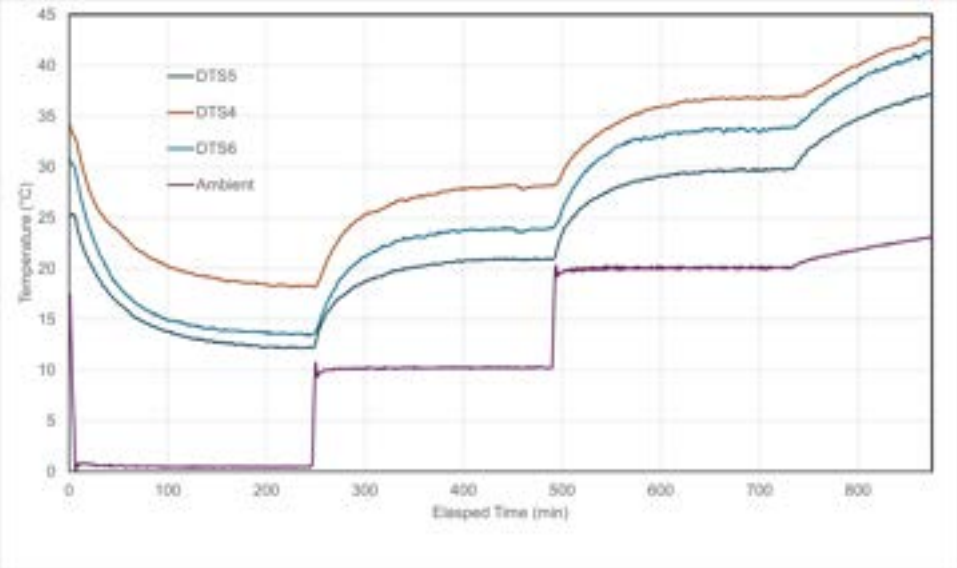

Figure A.27: Temperature test with TEC disabled

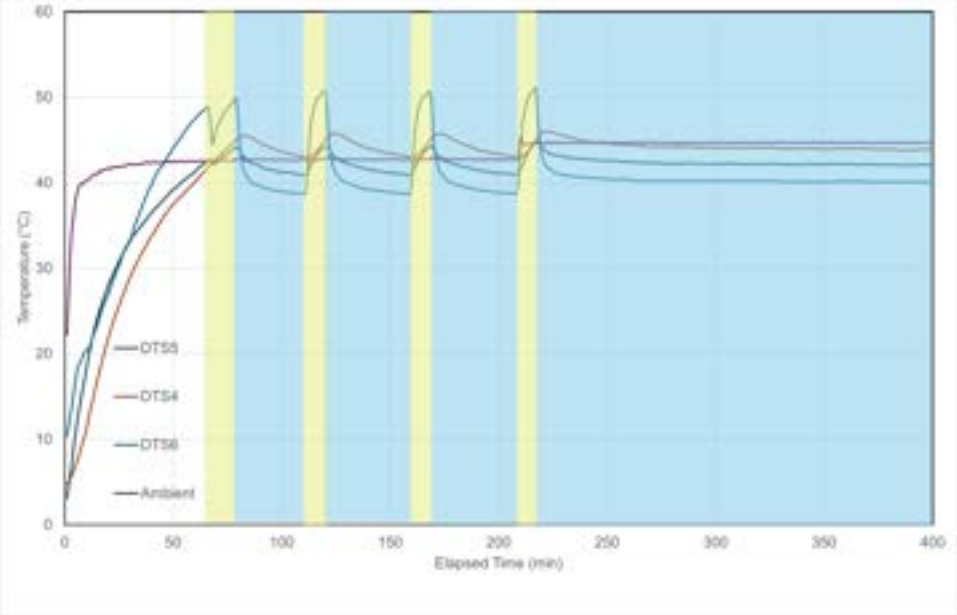

Figure A.28: High Temperature Test

controller setpoint to turn the TEC power on at 15° C instead of 8° C. Circulating the air inside the enclosure helps distribute the heat generated by the components.

A.3 Full-System RF Measurements

This section provides details on several RF measurements that were made on the fully integrated SEA sensor. The measurements encompassed the entire RF system as shown in Figure A.1 including the preselector and RSA, but excluding the antenna. The system also included a 30.5 m coaxial cable between the preslector and SPU, along with

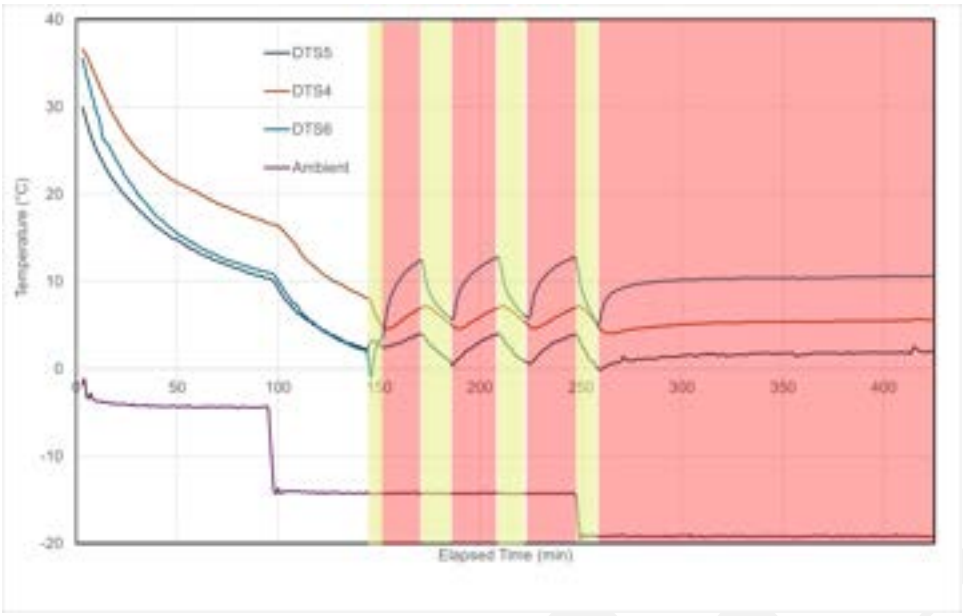

Figure A.29: High Temperature Test

a 6 dB attenuator in the SPU (component ATT). The combined insertion loss of the cable and attenuator was 15 dB. Note that while the length of the coaxial cable will likely differ for the deployed sensors due to the unique conditions at each site, the attenuator value will be changed as needed to achieve roughly the same combined insertion loss.

Further details on these measurements are discussed in the subsections that follow.

A.3.1 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 A.30.

This measurement shows the total system gain within the CBRS band, except for the upper 20 MHz, to be approximately 31.5 dB to 32.5 dB. At the upper end, from 3680 MHz to 3700 MHz, and gain falls from 31.5 dB to 15 dB. This gain characteristic closely follows the frequency response of the bandpass filter (BPF), as shown in Figure A.2. 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.

A.3.2 Noise Figure

The system noise figure is measured as part of the SCOS calibration action. This measurement follows the Y-factor technique and used the Keysight 346B noise source built into the preselector (designator NSR). With the noise source both enabled and disabled, IQ data is acquired from the RSA with a sample rate of 14 mega samples per second (Msps). From these IQ time series the mean power is computed, and subsequently the Y-factor and noise figure are calculated.

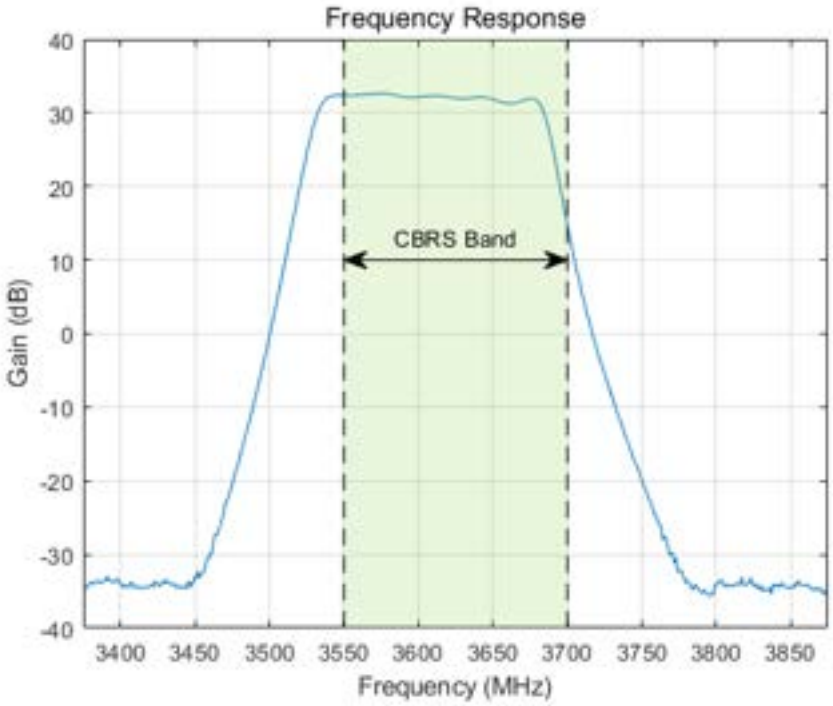

Figure A.30: Full System Frequency Response

This measurement is repeated at a total of 18 frequencies covering 15 measurement channels within the CBRS band and several channels just outside. The noise figure results for the first five Rev. 1 sensors is shown in Figure A.31.

The measured system noise figure is just under 4 dB for all measurement channels within the CBRS band, except for the upper two channels, Slight degradation is observed in the 3685 MHz channel, and further degradation of approximately 6 dB to 8 dB is observed in the 3695 MHz channel. This degradation is do to the roll-off of the bandpass filter at the upper end of the CBRS band, as shown in Figure A.2. Any increase in the insertion loss of the filter has a direct impact on noise figure.

A.3.3 Gain Compression and Sensor Overload

In addition to CW based compression testing described in Section A.1.1.4, the verification bench testing and acceptance (VBA) process was also used to evaluate the overload and compression of the sensors. The curves generated in the VBA process, as shown in Figure A.34, similarly trace out compression curves, except with a Gaussian noise source instead of a CW tone. The Gaussian noise source has a non-zero Peak-to-Average Power Ratio (PAPR), so the different effects of sensor overload (as indicated by the red markings on the curves) and sensor compression (indicated by the rolloff of the curve from the linear trend) can be seen independently. Because of the behavior of the overload flag in the signal analyzer, it is triggered well before actual compression occurs, since it will trigger on the peak measured value in the signal analyzer.

A.3.4 Calibration

In-field calibrations are performed on a periodic basis using an included noise diode in the Preselector to maintain calibration on the sensor over time. These noise diodes are first calibrated in the National Institute of Standards and Technology (NIST) radiometer for each 10 MHz channel and the resultant average excess noise ratio (ENR) is

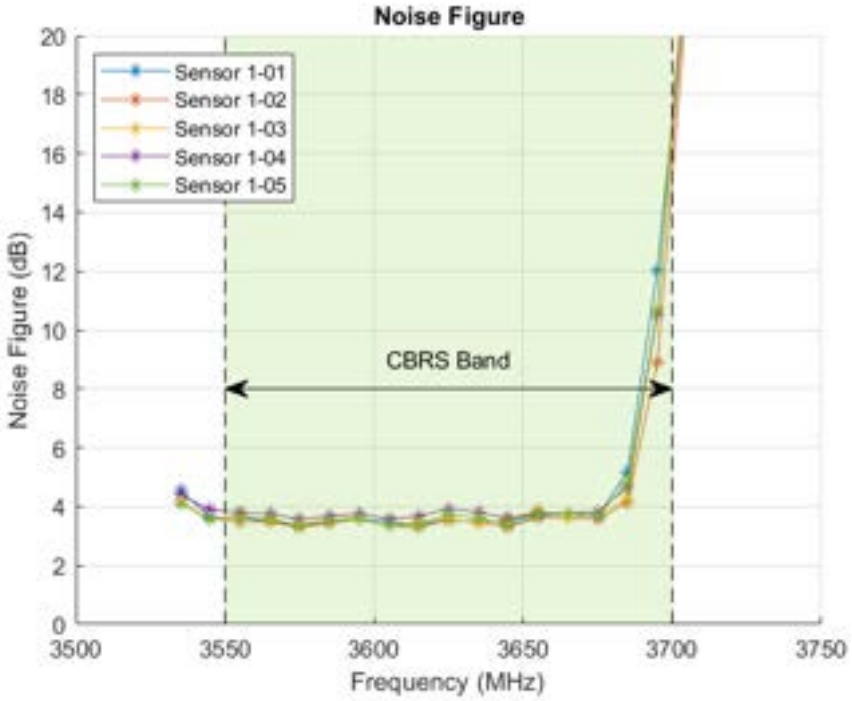

Figure A.31: Full System Noise Figure

configured on a sensor-by-sensor basis in the sensor configuration. The noise diode is then used to perform a Y-factor calibration of the sensor according to the equations below.

$$y = \frac{P_{\text{on}}}{P_{\text{off}}} \tag{A.1}$$

$$NF = \frac{\text{ENR}}{y - 1} \tag{A.2}$$

$$G_{\text{eff}} = \frac{P_{\text{on}}}{k_B T B(\text{ENR} + \text{NF})} \tag{A.3}$$

These calibrations are at a reference plane at the face of the noise diode in the preselector. In order to have the calibration at the reference plane of the antenna terminal an additional correction is required. This correction factor is calculated during the VBA procedure described Section A.4 and initially loaded onto the sensors during installation. The translation only goes through passive components in the sensor, so it is not expected to change over time during sensor operation, hence the calibration at the antenna reference plane is only performed once at installation.

Calibrations are performed for each setting of the signal analyzer that affect RF performance. For example, each frequency, reference level, attenuation, and use of hybrid sampling are calibrated separately. The calibration factors are then matched to the acquisition settings and front-end gain is corrected for in the collected data by scaling the IQ samples directly off the signal analyzer prior to filtering with the infinite impulse-response (IIR) filter. An example characterization of the extracted calibration values is shown in Figure A.32. Each of the measured channels can be successfully calibrated using the on-board noise diode.

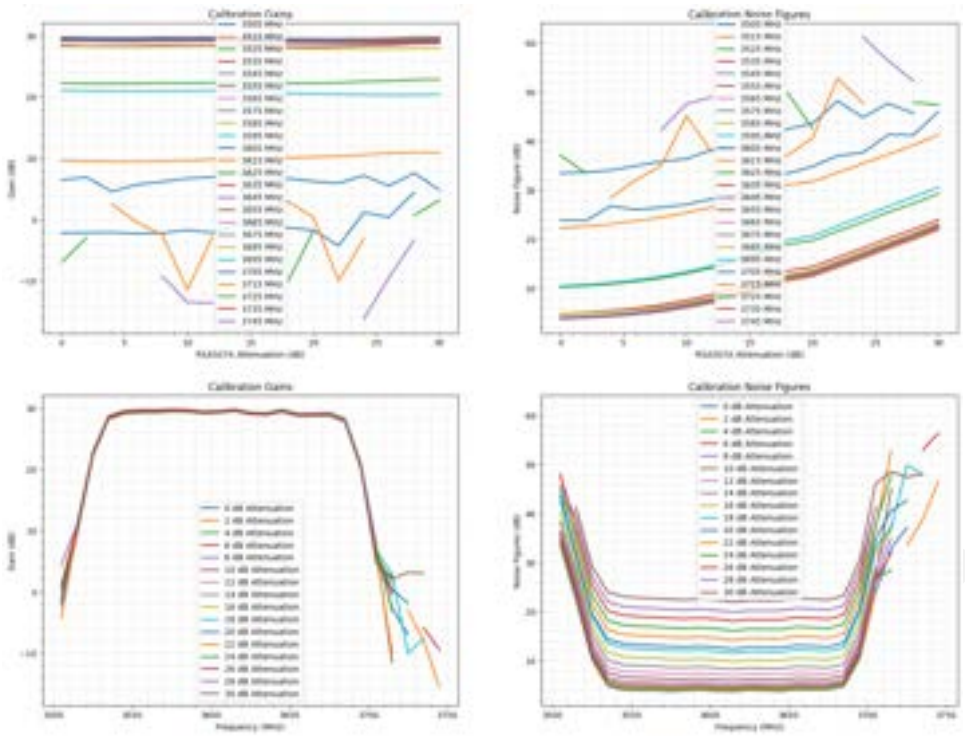

Figure A.32: Example full system calibration from 3500-3750 MHz. The measurement band, 3530-3710 MHz can be reliably calibrated using the on-board noise diode.

A.4 Sensor Verification and Acceptance Tests

The sensor verification tests are similar to those described in the previous NASCTN CBRS SEA Always-On DPA Test Plan [2, Chapter 7]. Initial tests are performed in the laboratory, post-construction, as a quality control procedure. Once a sensor's deployment location and site requirements are defined, they are fitted with appropriate RF cables and the VBA process is repeated to account for the new cables. These results are saved for verification during installation.

Upon installation, the VBA process is performed a final time on the as-installed sensor. Once the sensor SPU and preselector are mounted and the cable assembly is run between the two, the VBA process is performed to account for any differences caused by the new RF environment. During this, quality and installation verification is ensured by comparison of field VBA and pre-deployment results. The results of this final VBA procedure are then used to create correction values to account for the translation of the calibration reference plane to the antenna terminal.

A.4.1 RF Characterization

The RF verification and calibration follow the previous test plan [2, Section 7.3.1] and an example of this procedure follows. It is performed primarily using a a programmable noise probe designed to have an adjustable ENR. The probe is calibrated in the NIST radiometer for several attenuator settings to give a precise and accurate ENR for calibration (see Figure A.33), thus the sensor calibration is traceable to the NIST radiometer.

To perform the offset calibration, the noise probe is attached to the reference plane of interest, in this case the preselector antenna input, connected to a controller, and the attenuator is set to the maximum attenuation to give a measurement of thermal noise with a matched termination. Through the SCOS software, an onboard calibration is run to reference the calibration plane to the on-board noise diode. Data product actions are then taken at the maximum noise probe

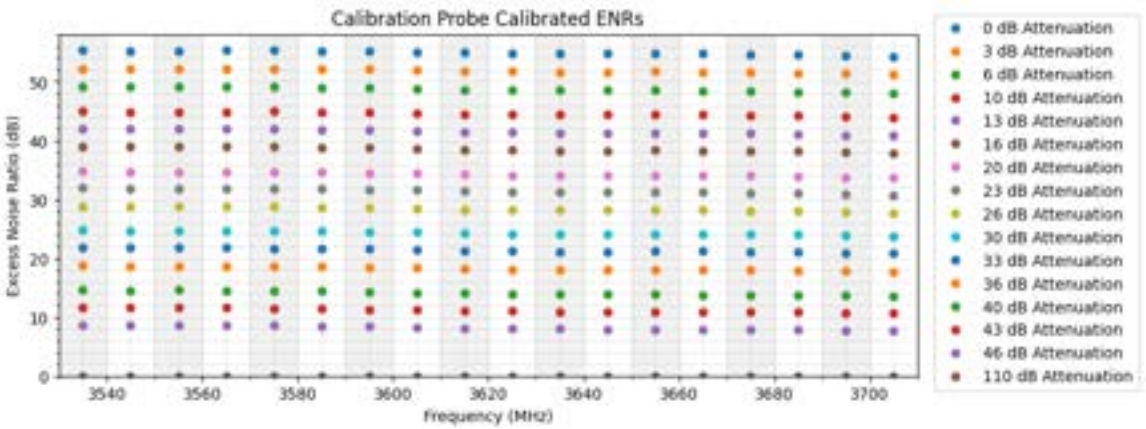

Figure A.33: Measured ENR for Programmable Noise Probe

attenuation setting and then all calibrated attenuation settings (note, these may be down selected after analysis of the data if not all settings are required).

At each frequency, the mean channel powers are extracted by averaging the Power versus Time data product for each attenuation setting. In Fig. A.34, these powers are then related according to the Y-factor equations A.1, A.2, A.3, and plotted versus the ENR of the probe attenuation they were taken with. The linear region (i.e. uncompressed region) is determined and fit to a linear model used to extract the noise figure of the system. The same linear range is then used to plot the linear channel power with the noise temperature of the calibration probe at the attenuation setting it was taken at. A linear fit to this trend is used to extract the gain of the system. It should be noted that the extracted noise figure of the system is of the entire system while the extracted gain is the calculated power offset from the on-board noise diode calibration plane to that of the calibration probe. This process can be seen in Figure A.34 for each measured channel.

A series of checks are then performed to ensure that the calibration was successful. First, the calculated gain values are corrected using the measured channel powers to offset the gain between the two calibration planes. These powers are then plotted against the thermal noise powers with the noise temperatures corresponding to the calibration probe attenuation settings. These must closely match to ensure that the appropriate offset was calculated. Next, the noise figure is translated from the on-board noise diode calibration plane to the calibration probe. Since all components along the path are passive, this should also scale linearly with the gain between them. It is also ensured that these values scale appropriately to confirm that the previous fits were successful.

Next, the cables are deembeded and extended. Any cable/adapter assembly required to connect the calibration probe to the sensor are measured for S21. The S21 value at the center of each 10 MHz channel is then removed to transition the calibration plane to the sensor port. The cable connecting the preselector to the antenna is then measured for its S21 and is similarly used to extend the calibration plane to the antenna port. These final offset gains are then put into a calibration file unique to the sensor and loaded into the SCOS software where it is applied in a similar manner to the on-board calibration offsets.

Finally, the S11 parameters are measured using a portable network analyzer in the field for the as-installed sensor and antenna. These will be packaged with the data and will allow for the final correction based on the mismatch at this plane and will include any effects of the antenna pattern reshaping due to nearby objects.

A.4.2 Variability Analysis

To check the variability of the above analysis, multiple data points were taken at each calibration probe attenuation of 20 dB, 30 dB, 40 dB, and 110 dB to give the linear trend as well as the matched "0 dB ENR". The above process was then carried out for all permutations of these data points to indicate the variability in the VBA calibration process. In Fig. A.35, the results of fitting each permutation are shown.

With nearly 45,000 iterations, the total spread of offset gains was on the order of 0.1 dB and the noise figures were similar except for the higher frequency channels where the filter attenuation increased the system noise figure. In the linear fits for these channels shown in Figure A.34, it can be seen that there is more non-linearity in the fitted region, hence a higher uncertainty is expected. This process is actively being iterated upon and refined and may change the attenuations used in the calibration to improve the variability in the fitting process.

A.4.3 System and Software Verification

In the VBA checkout and during installation, the system and software are ensured to be functioning. This involves a series of steps that exercise all components within the system and software, identifying that each is functioning as intended:

1. Assemble sensor and connect 50 ohm load to preselector antenna port
2. Power on the sensor and allow it to come to thermal equilibrium
3. Ensure the preselector light is illuminated
4. Using a computer connected to a different modem on the same LTE network:
5. Connect to the preselector through the web UI
6. Toggle each relay and ensure it updates with audible clicks if applicable
7. Connect to the webswitch through the webUI
8. Toggle each relay and ensure it updates with audible clicks if applicable
9. Ensure that the processor and the signal analyzer both power back up after toggling their switch port
10. Connect to the Next Unit of Computing (NUC) through secure shell (SSH) to ensure remote administration
11. Connect to the SCOS web interface and log in
12. Connect to the modem web configuration and log in to ensure it can be remotely administered
13. Once the UPS is charged, remove alternating current (AC) power and ensure that the sensor remains powered, then restore AC power
14. Connect to the processor directly through the SPU ethernet port:
15. Run a SCOS calibration action and ensure the output calibration values are expected given the calibrated ENRs of the internal noise diode
16. Run a SCOS data product action and ensure the measured noise values match kTB + NF for each channel
17. Connect the antenna cable and antenna
18. Run a SCOS data product action and ensure that the expected signal environment is observed

These steps are performed in addition to the RF calibration above, and are executed in the laboratory, as a pre-deployment check, and finally in the field during and post-installation. Once the sensor is confirmed to be functioning correctly, it is tagged and recorded to ensure that each component is accounted for. Finally, the sensor is configured to run calibration actions at a regular interval and data product actions in quick succession.

Post-installation, the initial data from the sensor is analyzed for each channel to investigate if any of the dynamic range adjustments need to be made. The sensor's configuration is then updated to include the new settings, both in the calibration action as well as in the data product actions. This process is repeated until regular overloads cease. This process is regularly re-checked to ensure that new environmental variables do not warrant changes to these dynamic range settings.

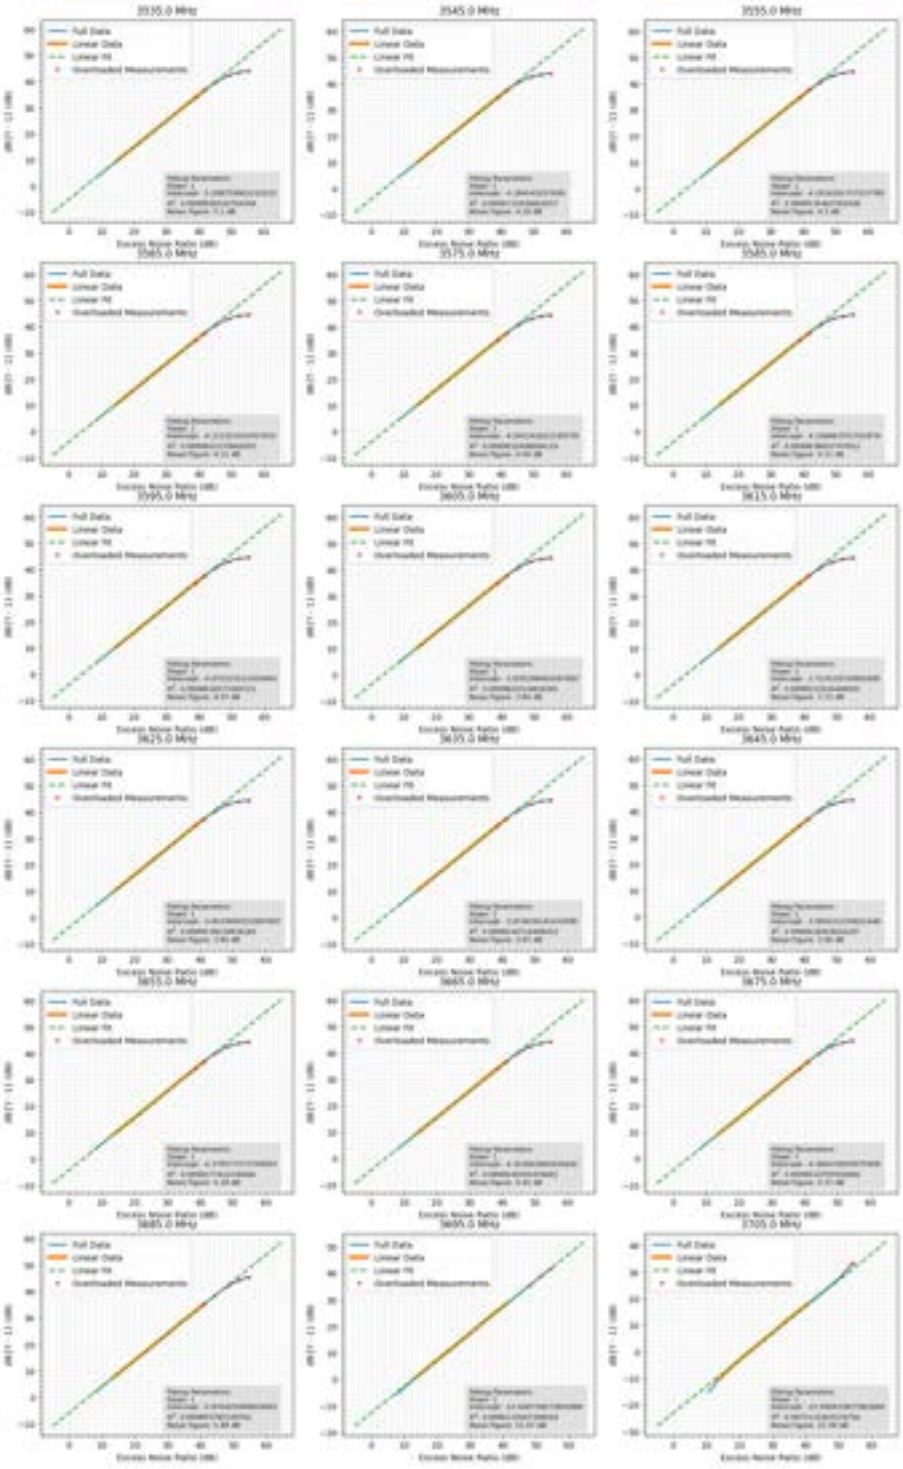

Figure A.34: Programmable Noise Source Calibration with Linear Fit for Each Channel

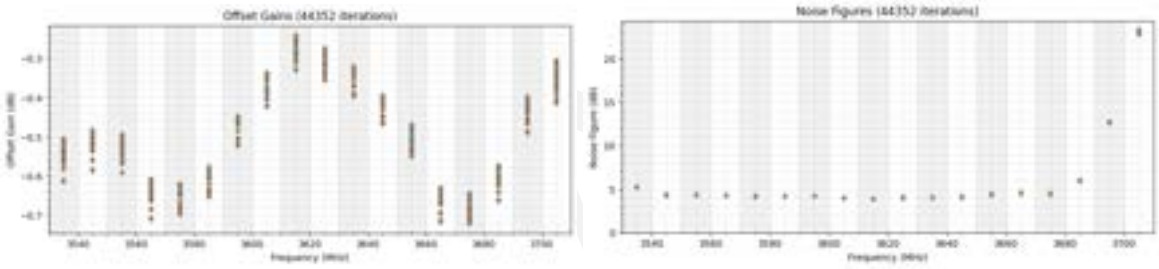

Figure A.35: Variability of gain and noise figures for 45352 iterations of four point fits

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] Duncan A. McGillivray and Todd Schumann and Thao T. Nguyen and Douglas Boulware and M. Keith Forsyth and Dazhen Gu and C. Lee Joyce and Elyssa Kaplan and Mark and Daniel G. Kuester and Joseph R. Mruk and Anthony W. Romaniello and Aric W. Sanders and Robert Scheeler and Stephen Segro, "DRAFT CBRS Ecosystem Sharing Assessment (SEA) Test and Metrology Test Plan: Characterize aggregate emissions within the CBRS band 3550-3700 MHz in at least one Always-On DPA ," 2023. [Online]. Available: https://www.nist.gov/system/files/documents/2023/12/19/NASCTN%20CBRS%20SEA%20Task%202%20Draft%20Test%20Plan%20-%20December%202023.pdf (Last accessed: 15 February 2024).

[3] D. Boulware, A. Romaniello, R. Dorch, and M. Cotton, "An Analysis of Aggregate CBRS SAS Data from April 2021 to January 2023," National Telecommunications and Information Administration, Tech. Rep. TR-23-567, 2023. [Online]. Available: https://www.ntia.gov/sites/default/files/publications/tr-23-567.pdf

[4] 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: 08 Jan 2024).

[5] 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).

[6] 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).

[7] CBRS WInnForum Standards, "Requirements for Commercial Operation in the U.S. 3550–3700 MHz Citizens Broadband Radio Service Band, Document WINNF-TS-0112, Version V1.9.1," 2020. [Online]. Available: https://cbrs.wirelessinnovation.org/release-1-standards-specifications (Last accessed: 08 February 2024).

[8] Wireless Innovation Forum, "SAS Testing and Interoperability Repository." [Online]. Available: https://github.com/Wireless-Innovation-Forum/Spectrum-Access-System (Last accessed: 08 February 2024).

[9] ITU, "Recommendation ITU-R P.2108-1 - Prediction of clutter loss." [Online]. Available: https://www.itu.int/rec/R-REC-P.2108-1-202109-I/en (Last accessed: 08 February 2024).

[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