An End-to-End Radio Frequency Interference Monitoring System

by

Nicholas Bruce
B.Eng., University of Victoria, 2016
M.ASc., University of Victoria, 2018

A Dissertation Submitted in Partial Fulfillment of the
Requirements for the Degree of

DOCTOR OF PHILOSOPHY

in the Department of Electrical and Computer Engineering

We acknowledge and respect the Ləkʷəŋən (Songhees and Xʷsepsəm/Esquimalt) Peoples on whose territory the university stands, and the Ləkʷəŋən and W̱SÁNEĆ Peoples whose historical relationships with the land continue to this day.

An End-to-End Radio Frequency Interference Monitoring System

by

Nicholas Bruce
B.Eng., University of Victoria, 2016
M.ASc., University of Victoria, 2018

Supervisory Committee

Dr. P. F. Driessen, Supervisor
(Department of Electrical and Computer Engineering)

Dr. S. Harrison, Supervisor
(Department of Electrical and Computer Engineering)

Dr. T. Robishaw, Departmental Member
(Department of Electrical and Computer Engineering)

ABSTRACT

This dissertation presents the design, implementation, and philosophy of an end-to-end radio frequency interference monitoring system designed primarily for radio astronomy observatories. Motivated by the growing density of anthropogenic emitters and the corresponding threat to ground-based radio astronomy, the work describes a modular monitoring pipeline that includes hardware, signal processing, and machine learning. This work introduces several novelties: a gain-stable wideband receiver providing calibrated long-term radio frequency environment comparisons, a direction-of-arrival estimation method allowing use of arbitrarily arranged antennas, a digital down-converter architecture enabling insight into individual signal detections, and an unsupervised learning framework to perform modulation recognition and novelty detection. The dissertation also emphasizes the importance of modular, software-defined architectures and outlines future work towards real-time algorithms for next-generation interference management.

Contents

List of Tables

List of Figures

List of Acronyms

ADC	analog-to-digital converter
AI	artificial intelligence
ARTTA	Advanced Radio Telescope Timing Array
CGEM	Canadian Galactic Emission Mapper
CHIME	Canadian Hydrogen Intensity Mapping Experiment
CHORD	Canadian Hydrogen Observatory and Radio-transient Detector
CW	continuous waveform
DDC	digital down-converter
DFT	discrete Fourier transform
DoA	direction-of-arrival
DPDK	Data Plane Development Kit
DRAO	Dominion Radio Astrophysical Observatory
DSP	digital signal processing
DVA	Dish Verification Antenna
EMC	Electromagnetic Compatibility
ENOB	effective number of bits
ESPRIT	Estimation of Signal Parameters via Rotational Invariance Technique
FIFO	first in, first out

FM	frequency modulation
FPGA	field-programmable gate array
FSPL	free space path loss
FT	Fourier transform
FFT	fast Fourier transform
FWHM	full width at half maximum
GPIO	general purpose input/output
GPU	graphics processing unit
IF	intermediate frequency
IP	Internet Protocol
IP1dB	input power 1 dB compression point
ISED	Innovation, Science and Economic Development Canada
ITU	International Telecommunication Union
ITU-R	International Telecommunication Union Radiocommunication Sector
JAG	John A. Galt telescope
LLM	large language model
LNA	low-noise amplifier
LSL	LWA Software Library
LWA	Long Wavelength Array
LWA-SV	Long Wavelength Array Sevilleta Station
ML	machine learning
MUSIC	Multiple Signal Classification
NFA	noise-figure analyzer

OP1dB output power 1 dB compression point

OSPFB oversampled polyphase filterbank

PFB polyphase filterbank

PSD power spectral density

QSFP quad small form-factor pluggable

RA radio astronomy

RF radio frequency

RFI radio frequency interference

SDR software-defined radio

SigMF Signal Metadata Format

SFDR spurious-free dynamic range

SFM Solar Flux Monitor

SKA Square Kilometer Array

SNR signal-to-noise ratio

STFT short-time Fourier transform

TBN Transient Buffer Narrow

TEA thermoelectric assembly

TTL transistor-transistor logic

UDP User Datagram Protocol

USRP Universal Software Radio Peripheral

VITA VMEbus International Trade Association

VNA vector network analyzer

ZeroMQ Zero Message Queue

ACKNOWLEDGMENTS

I would like to thank:

Emily for putting up with this

Peter for always making me think about the bigger picture

Stephen for answering an impossible number of questions via SMS

Tim for telling me I was going to get my PhD

The folks at Herzberg Astronomy and Astrophysics for support in finishing my degree

The document was typeset in LaTeX across a variety of Linux distros, Overleaf.com, and macOS. Zotero was used for citation management and Git was used for version control.

The majority of simulation and data analysis was done using Python 3. The analyses that needed more compute than a standard computer could offer were run on Digital Research Alliance of Canada resources (formerly ComputeCanada).

Plots were made using a combination of the Matplotlib and Plotly Python packages, as well as directly in LaTeXusing the Pgfplots package. Diagrams and drawings were done in LaTeX using PGF/TikZ.

> The only absolute knowledge attainable by man is that life is meaningless.
>
> Leo Tolstoy, A Confession

Chapter 1

Introduction

You're an *interfering* engineer?

Helen Bruce, 2024

1.1 Problem Description

Monitoring the radio frequency (RF) spectrum is an increasingly difficult and important task. Many disparate groups undertake monitoring efforts including governments, telecommunications companies, militaries, and science centers such as those hosting space missions and physics research. Of these groups, science centers hosting ground-based astronomical radio telescopes and other radio-quiet instruments are unique in that every anthropogenic radio signal is, by definition, radio frequency interference (RFI)[1][2]. Ideally, a radio-quiet site would be free of all anthropogenic radio signals because even weak emitters can corrupt the far weaker astronomical radio signals. In an effort to avoid interference sources altogether, some groups have recently conducted concept studies for lunar-farside radio astronomy (RA) [1, 2]. However, RFI will remain an issue for ground-based observations.

There has recently been substantial growth in spectral occupancy, driven by the increasing ubiquity of wireless communications and the development of satellite mega-constellations such as SpaceX's Starlink and Amazon's Leo (previously Kuiper). Individual satellites in these constellations operate across bandwidths up to 250 MHz, and the

[1]For the purposes of this dissertation, "spectrum monitoring" is referred to as RFI monitoring, as all anthropogenic signals in this context constitute interference.

[2]A notable exception is radar astronomy, where the telescope functions as both transmitter and receiver for high-resolution imaging.

large number of satellites in these constellations will ensure near-continuous sky coverage [3]. Recent studies have found unintentional satellite emissions well outside licensed frequency allocations [4]. The aggregate effect of these unintended emissions has the potential to end ground-based astronomy as the number of satellites overhead grows into the hundreds of thousands or even millions [5, 6].

Terrestrial RFI sources can be categorized by whether emissions are intentional and by the regulatory framework governing their use.

Intentional emitters can be further divided into licensed and unlicensed transmitters. Licensed transmitters can legally operate in specific frequency bands and include high-power sources intended to operate at long ranges such as commercial radio and television stations, telecommunications towers, and airport radar. Unlicensed transmitters operate in open bands and can use short-range technologies such as Wi-Fi and Bluetooth.

In unintentional emitters, RFI is a side effect of normal electrical operation; examples include vehicle ignition systems, high-voltage power lines, and many consumer electronics.

Modern devices often fall into more than one of these categories. For example, cellular telephones are simultaneously licensed intentional emitters (LTE/5G), unlicensed intentional emitters (Wi-Fi/Bluetooth), and unintentional emitters (internal electronics).

The RF spectrum, occupied by both satellite and terrestrial sources, is an increasingly valuable and contested resource that demands robust monitoring solutions.

I have used two RA observatories as field sites for the RFI monitoring solutions that I present in this dissertation. In particular, my research is conducted primarily at Dominion Radio Astrophysical Observatory (DRAO) in British Columbia, Canada, with additional work at the Long Wavelength Array (LWA) radio astronomy telescope in New Mexico, USA.

A typical approach to RA observatory RFI monitoring is to build an analog receiver and connect it either to a commercial spectrum analyzer or to an existing telescope spectrometer. In both cases, it is trivial to generate large volumes of data and non-trivial to make use of those data. When the system is not well calibrated, observing long-term trends in the RF environment is impossible. When spatial localization is attempted, it is typically achieved using steerable RA telescopes (at the cost of losing astronomical observation time), or a dedicated steerable RFI antenna (at the cost of not continuously observing the entire RA facility). To the best of my knowledge, interferometric synthesis-imaging techniques have not previously been applied to the design of an RFI monitoring system.

While high-fidelity feature extraction methods such as cyclostationary signal processing offer a robust framework for environment analysis [7], their application to real-time RFI monitoring is precluded by their extreme computational complexity. This mirrors challenges in computational acoustic scene analysis, in which characterizing multiple overlapping sources necessitates more efficient algorithms [8].

The problem that I address in this work is the absence of a satisfactory end-to-end RFI monitoring system—one capable of generating calibrated data suitable for both long-term trend analysis and transient detection, spatially localizing detected signals, and extracting actionable information from those data, all in real time.

1.2 Contributions

In the remainder of this dissertation, I cover several components of a robust end-to-end RFI monitoring system:

- the design and implementation of a gain-stable RFI monitor;
- a method for estimating signal direction-of-arrival (DoA) with arbitrarily arranged receiver antennas;
- a method for synthesizing channelized signals into a baseband representation; and
- a method for categorizing detected signals without prior labelled training data.

1.2.1 Wideband Gain-Stable Receiver

This work began with the development, primarily by Stephen Harrison, of a prototype RFI monitor at the DRAO [9]. It operated for approximately five years before I undertook its redesign as a production instrument. This required establishing requirements, redesigning the monitor where necessary, and experimentally verifying compliance with those requirements.

While establishing the requirements, I found that the standard method for defining gain drift in a radiometer does not account for the influence of integration. This omission motivated the development of a novel method for characterizing the effect of gain drift on integrated data.

This work also introduces the concept of using the RFI monitor as a platform for RF environment analysis both over the long term, by stabilizing and calibrating the monitor properly, and in the short term, by detecting transients and anomalous signals in real time and responding to them. I argue that storing large quantities of spectral data for future reanalysis is the wrong approach. A better approach is to develop forward-looking systems and algorithms, extend them as new needs arise, and treat selective data loss as an acceptable cost of real-time operation.

1.2.2 Signal DoA Estimation

This idea was initially motivated by my prior work on ionospheric imaging [10, 11]. I wanted to extend this work to a higher-resolution array with more elements such as the LWA but needed a method to efficiently obtain direction of arrival. As I began to work with both the LWA and the DRAO RFI monitoring group, it became apparent to me that DoA is often a secondary consideration for RFI monitoring. Most RFI monitors described in the literature are either omnidirectional or mechanically or electronically steered. All of these scenarios result in either no DoA information or an inability to survey the entire local RF environment at any given time.

A series of omnidirectional antennas used in concert could continuously monitor the entire local RF environment while simultaneously providing DoA information.

From this, I developed a visibility modelling-based method to determine the elevation and azimuth angles for a signal received at a quasi-randomly arranged (2-D) array [12, 13]. The quasi-random arrangement of antennas was an important consideration since at many sites antenna locations are restricted by existing infrastructure and geography. The method achieves the maximum angular resolution available from a given number of antennas when no two antenna baselines are repeated. Simulations show that this method has comparable accuracy to subspace-based DoA algorithms such as Multiple Signal Classification (MUSIC) [14] and measured data show that it is both more accurate and faster than other visibility-based methods such as RA all-sky imaging [15].

1.2.3 Signal Down-Conversion

Because the guiding principle is that data should be harvested in real time rather than stored indefinitely, discussions with my supervisors, Peter Driessen and Stephen Harrison, raised the question of what data should be collected and how. Because the primary data product from most monitoring systems is integrated spectral data, some information can be directly extracted, such as bandwidth, duration, time of day, and power, but the complex baseband time series of each signal would yield deeper insights. From such a representation we could expect to determine complex modulation schemes, demodulate messages, and choose to treat that time series arbitrarily, including rechannelizing to a different bin count.

Consider also the issue of storing data. As mentioned previously, in a monitoring system where the observed bandwidth is very high, it is at best impractical and at worst impossible to store a continuous stream of data to disk. The value per storage unit is low and further analysis will require re-processing a massive dataset. Alternatively, consider a system where only the message bandwidth of each signal is stored or, more selectively, only metadata describing each signal are stored. This shift from a continuous, undifferentiated data stream to a structured collection of high-value signal records would substantially reduce the storage requirement and simplify downstream analysis.

There are several well-documented digital down-converter (DDC) architectures [16] and commercial products available that I expand upon in the dissertation. Largely, they focus on communications applications where one signal at a time is down-converted using the classic or naive method. The naive method of down-converting a signal to its complex baseband equivalent is to frequency shift the entire spectrum until the signal is centered around 0 Hz, use a signal-width low-pass filter to reject the rest of the spectrum, and then downsample to reduce the output sample rate.

This architecture requires that each signal be passed through the entire DDC chain serially. During each DDC process, the entire wideband time series has to be multiplied by a complex exponential and the product stored and filtered before downsampling, while the initial wideband time series must remain untouched for use with the next signal as illustrated in fig. 1.1. Many optimized DDC implementations are built around field-programmable gate array (FPGA) platforms [17, 18]. Despite these optimizations, they still use the naive down-conversion method and are therefore unsuitable for real-time monitoring of a wideband environment in which hundreds or thousands of signals appear simultaneously and a snippet of each must be down-converted.

One method to handle this was introduced in [19–21], in which a twice-oversampled analysis filterbank channelizes the input time series and then a twice-undersampled synthesis filterbank recombines the channels containing signal fragments. This allows the entire spectrum to be prepared for down-conversion without requiring each signal to be shifted independently, and the signals are already decomposed into baseband fragments ready for reconstruction.

More recently, a publicly available but not formally published paper from a signal processing workshop described a similar implementation in GNU Radio [22]. While their choice of filter does not allow perfect channel synthesis, the paper shows that there is

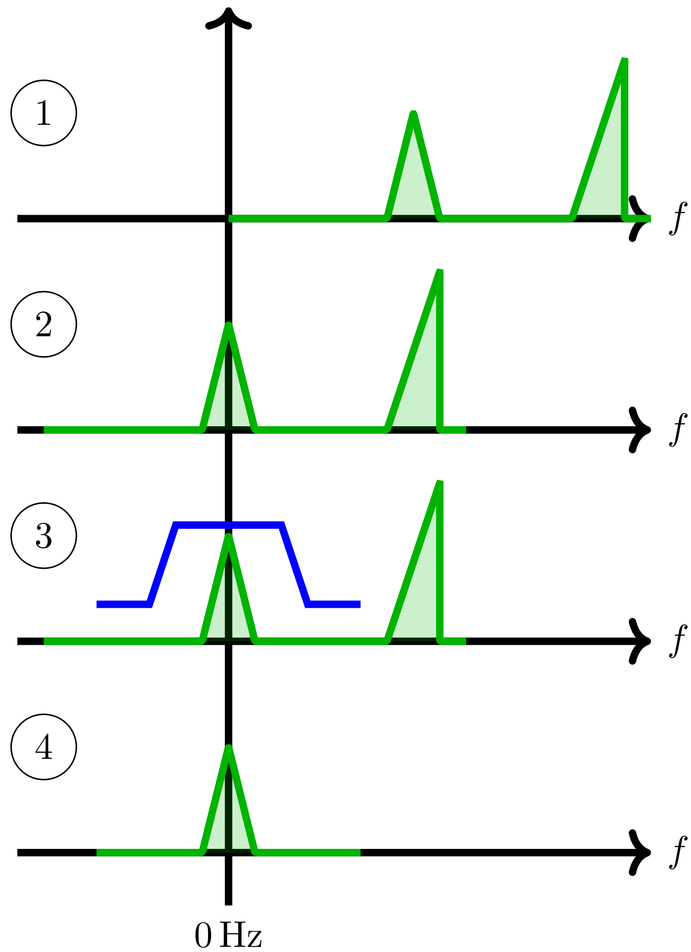

Figure 1.1: A naive DDC implementation. A wideband time series containing multiple signals (1) is shifted (2) and filtered to isolate a signal (3), then downsampled (4) to obtain the isolated signal's baseband time series. The process is repeated for each signal of interest.

general interest from the wider software-defined radio (SDR) and digital signal processing (DSP) communities in improved signal down-conversion methods.

While the method was introduced in [19–21] and later implemented in [22], in no case are the implementation details fully described, the computational efficiency rigorously assessed, or an end-to-end system presented. In this dissertation, I revisit these methods and extend them by introducing a fractionally oversampled architecture.

1.2.4 Unsupervised Signal Analysis

The next stage of an end-to-end RFI monitoring system is to process the complex baseband time series discussed above. Converting every detected signal to a complex baseband time series is complicated by the sheer volume of signals—hundreds of thousands per day across the 0–2 GHz bandwidth continuously observed by the RFI monitor at DRAO. Manual categorization at this volume is impractical, and many categorization methods require at least a partially labeled dataset—labels that demand significant expert effort to produce.

Given the substantial resources invested in machine learning (ML) by both industry and academia, and the rapid progress of the field, I believe that methods developed for general-purpose ML are ready for transfer to the RFI monitoring domain. Even without specializing in ML, one can transfer lessons from the field to RFI monitoring. One particularly influential paper introduced SimCLR [23], a framework for unsupervised image clustering. In SimCLR, augmented views of the same image are pulled together in representation space while views from different images are pushed apart, allowing the model to learn image similarity without labels. This raised the question of whether a similar approach would generalize to time-domain data. This line of questioning produced promising results, though much remains to be done. I introduce SigCLR, a method capable of distinguishing between signal types and generalizing across multiple domains.

The primary application is modulation recognition, enriching the metadata associated with each signal detection. SigCLR can also function as a novelty detector, flagging signals that bear no resemblance to any previously observed class. Unlike traditional ML models that are limited to a fixed set of classes, this unsupervised approach identifies novelty through the clustering structure of the learned representations. Because known signals group together, any new or anomalous RFI will fail to join these clusters, and can therefore be flagged as a novel detection.

1.2.5 Summary of Contributions

In summary, the main contributions of this dissertation are:

- An improved method for characterizing the effect of gain drift on integrated radiometer data.
- A method for estimating DoA from arbitrarily arranged array geometries.

- A novel application of oversampled synthesis filterbanks to the computationally efficient reconstruction of detected interference signals.
- An unsupervised image categorization technique transferred to complex time-series radio data for modulation recognition and novelty detection.

1.3 Organization of the Dissertation

The remainder of this dissertation is organized as follows:

- Chapter 2 presents the requirements specification, design, and commissioning of a wideband RFI monitor.
- Chapter 3 introduces a novel DoA method for arbitrarily arranged array geometries.
- Chapter 4 develops a method for reconstructing the baseband time series of detected signals using twice-oversampled filterbank channels.
- Chapter 5 describes the application of an unsupervised ML technique, normally applied to image data, to complex time-series data.
- Chapter 6 presents conclusions and directions for future work.

As detailed in section 1.4, several of these chapters have been published or submitted for publication. To preserve the structure of each publication as incorporated into this dissertation, each chapter has its own set of appendices and a dedicated bibliography. The writing style also varies by chapter, reflecting the different conventions of engineering and astronomy publications.

1.4 Publications

At the time of this writing, the following chapters of this dissertation have been published:

- Chapter 2 has been published as [24].
- Chapter 3 has been published as [25].
- An earlier version of Chapter 4 has been published as [26].
- An earlier version of Chapter 5 has been published as [27].

References

[1] J. O. Burns, R. MacDowall, S. Bale, G. Hallinan, N. Bassett, and A. Hegedus, "Low Radio Frequency Observations from the Moon Enabled by NASA Landed Payload Missions," *The Planetary Science Journal*, vol. 2, p. 44, Apr. 2021.

[2] J. O. Burns, "Transformative science from the lunar farside: observations of the dark ages and exoplanetary systems at low radio frequencies," *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences*, vol. 379, p. 20190564, Jan. 2021.

[3] I. del Portillo, B. G. Cameron, and E. F. Crawley, "A technical comparison of three low earth orbit satellite constellation systems to provide global broadband," *Acta Astronautica*, vol. 159, pp. 123–135, June 2019.

[4] F. Di Vruno, B. Winkel, C. G. Bassa, G. I. G. Józsa, M. A. Brentjens, A. Jessner, and S. Garrington, "Unintended electromagnetic radiation from Starlink satellites detected with LOFAR between 110 and 188 MHz," *A&A*, vol. 676, p. A75, 2023.

[5] M. L. Rawls, H. B. Thiemann, V. Chemin, L. Walkowicz, M. W. Peel, and Y. G. Grange, "Satellite constellation internet affordability and need," *Research Notes of the AAS*, vol. 4, p. 189, Oct. 2020. arXiv: 2011.05168.

[6] Space Exploration Holdings, LLC, "Application for authority to launch and operate the SpaceX orbital data center system." FCC ICFS File No. SAT-LOA-20260108-00016, Jan. 2026. Narrative at 1-5. [Online]. Available: `https://fccprod.servicenowservices.com/icfs?id=ibfs_application_summary&number=SAT-LOA-20260108-00016`.

[7] W. Gardner, *Cyclostationarity in Communications and Signal Processing*. Electrical engineering, communications and signal processing, IEEE Press, 1994.

[8] Y. Li, W. Cao, W. Xie, Q. Huang, W. Pang, and Q. He, "Low-complexity acoustic scene classification using data augmentation and lightweight ResNet," in *2022 16th IEEE International Conference on Signal Processing (ICSP)*, vol. 1, pp. 41–45, 2022.

[9] S. Harrison, A. Lakhani, J.-P. Lavigne, N. Bruce, T. Robishaw, R. Messing, C. Baard, M. A. Smith, B. Robert, and D. Del Rizzo, "RFI Monitor System Design for the Dominion Radio Astrophysical Observatory," in *2021 URSI GASS*, 2021.

[10] N. Bruce, S. Harrison, R. Herring, and P. Driessen, "Development of a phased-array ionospheric imaging system," in *2019 IEEE Pacific Rim Conference on Communications, Computers and Signal Processing (PACRIM)*, (Victoria, BC, Canada), pp. 1–8, IEEE, Aug. 2019.

[11] N. Bruce, *Development of a phased-array ionospheric imaging system.* PhD thesis, University of Victoria, 2019.

[12] A. R. Thompson, J. M. Moran, and G. W. Swenson, *Interferometry and synthesis in radio astronomy.* Springer, 3 ed., 2017. arXiv: 1011.1669v3 ISSN: 09646639.

[13] G. B. Taylor, C. L. Carilli, and R. A. Perley, *Synthesis Imaging in Radio Astronomy II*, vol. 180. Astronomical Society of the Pacific, 1999. Publication Title: Synthesis Imaging in Radio Astronomy II ISSN: 1050-3390.

[14] R. Schmidt, "Multiple emitter location and signal parameter estimation," *IEEE Transactions on Antennas and Propagation*, vol. 34, no. 3, pp. 276–280, 1986.

[15] S. Ellingson and G. Hampson, "Detection and localization of L-band satellites using an antenna array," in *IEEE Antennas and Propagation Society Symposium, 2004.*, vol. 1, pp. 177–180 Vol.1, 2004.

[16] A. Abinaya and M. Maheswari, "A survey of digital down converter architecture for next generation wireless applications," *IOP Conference Series: Materials Science and Engineering*, vol. 872, p. 012037, June 2020.

[17] P. Sikka, A. R. Asati, and C. Shekhar, "Power- and area-optimized high-level synthesis implementation of a digital down converter for software-defined radio applications," *Circuits, Systems, and Signal Processing*, vol. 40, pp. 2883–2894, June 2021.

[18] G.-l. Shao, X.-s. Chen, and L. Yang, "Research and implementation of a high performance parallel computing digital down converter on graphics processing unit: Parallel computing digital down converter on GPU," *Concurrency and Computation: Practice and Experience*, vol. 29, p. e4042, Apr. 2017.

[19] F. Harris, C. Dick, X. Chen, and E. Venosa, "Wideband 160-channel polyphase filter bank cable TV channeliser," *IET Signal Processing*, vol. 5, no. 3, p. 325, 2011.

[20] F. Harris, E. Venosa, X. Chen, and B. D. Rao, "Polyphase analysis filter bank down-converts unequal channel bandwidths with arbitrary center frequencies," *Analog Integrated Circuits and Signal Processing*, vol. 71, no. 3, pp. 481–494, 2012.

[21] E. Venosa, F. J. Harris, and F. A. Palmieri, *Software Radio: Sampling Rate Selection, Design and Synchronization*, vol. 53. Springer, 2013. Publication Title: Analog Circuits and Signal Processing Issue: 9.

[22] T. W. Rondeau, T. J. O'Shea, and V. Tech, "Designing analysis and synthesis filter-banks in GNU Radio," in *Karlsruhe Institute of Technology's Workshop on Software Radio (WSR)*, (Karlsruhe), p. 9, 2014.

[23] T. Chen, S. Kornblith, M. Norouzi, and G. Hinton, "A simple framework for contrastive learning of visual representations," June 2020. arXiv:2002.05709 [cs, stat].

[24] N. Bruce, C. Baard, S. Harrison, M. Islam, A. J. Otto, D. Lagoy, R. Messing, B. Robert, T. Robishaw, and P. F. Driessen, "Wideband RFI monitor requirements, design, and commissioning at DRAO," *Publications of the Astronomical Society of the Pacific*, vol. 138, p. 065005, June 2026.

[25] N. Bruce, J. Stacey, and P. F. Driessen, "Visibility domain direction-of-arrival optimization for arbitrary 2-D arrays," *IEEE Access*, pp. 1–1, 2022.

[26] N. Bruce, S. Harrison, and P. Driessen, "DDC Pool: Efficient down-conversion of signals for RFI monitoring," in *Proceedings for RFI 2022*, 2022.

[27] N. Bruce, B. Moa, S. Harrison, and P. F. Driessen, "SigCLR: A contrastive learning approach to unsupervised modulation recognition and novelty detection," in *Proceedings of Radio Frequency Interference Conference — PoS(RFI2024)*, vol. 471, p. 025, 2025.

Chapter 2

Wideband RFI Monitor Requirements, Design, and Commissioning at DRAO

> I feel my temperature rising. Help me, I'm flaming, I must be a hundred and nine.
>
> Elvis Presley, Burning Love

2.1 Introduction

To operate Canada's DRAO effectively, we must monitor the primary frequency bands of all operating telescopes on the site for local sources of RFI. This monitoring effort primarily aims to preserve the radio-quiet environment of the observatory site for current and future science instruments. The DRAO hosts leading-edge instruments owned by the Canadian government and external collaborators (see § 2.2.1), all of whom directly benefit from this RFI monitoring effort.

The monitor serves the science community by confirming interference events in observation data. It further enables study of the statistical nature of the local RF environment to establish the suitability of proposed experiments. To complete these feasibility assessments, accurate long-term statistics characterizing the DRAO RF environment are crucial. Instruments from external collaborators, such as the Canadian Hydrogen Intensity Mapping Experiment (CHIME), can provide insight into RF environment trends on site (fig. 2.1), but they monitor only their specific frequency ranges of interest, are not typically designed to identify RFI sources, lack shared noise-floor characteristics with other instruments, and have finite operational lifespans. Consequently, a dedicated instrument is valuable for identifying long-term trends on site. A standardized comparison of spectra across multiple years is possible only if the instrument is well calibrated to remove the effects of instabilities, yielding accurate power measurements.

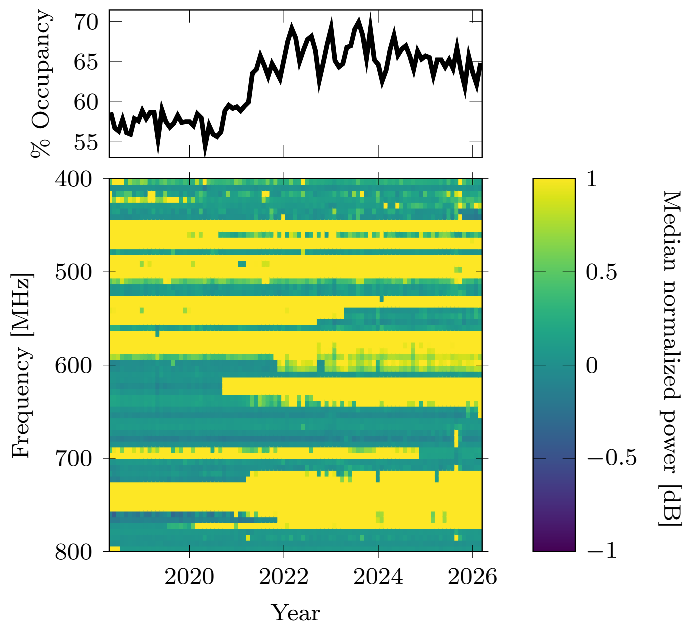

Figure 2.1: Annual snapshots of CHIME data normalized against a snapshot median (lower) showing an increasing trend in percent occupancy above 0 dB (upper) [1].

The monitor assists the local operations team in their daily effort to preserve the observatory's radio-quiet environment. Meanwhile, instrument builders rely on these calibrated power measurements to design hardware suited for typical on-site power levels [2]. Spectrum managers leverage the data to report accurate spectrum utilization and interference statistics to national and international regulators, including Innovation, Science and Economic Development Canada (ISED) and the International Telecommunication Union (ITU). Finally, the monitor system serves as a dedicated platform for researching advanced RFI detection and mitigation algorithms.

The monitor initially operated as a prototype for five years [3], during which time it streamed real-time data that detected various interfering sources, such as handheld cellular devices carried by delivery drivers, telematic data signals from vehicles parked on site, and an unshielded oscillator on a local telescope antenna. It also served as a foundational platform for related research, such as applying synthesis filterbanks to the problem of RFI identification [4] and performing unsupervised signal modulation recognition [5].

While the monitor operated as a prototype, telemetry data revealed that the analog signal chain—which must remain stable for accurate calibration—fluctuated with the outdoor temperature (fig. 2.2) and that the heater-cooler unit inside the RF enclosure was operating at maximum capacity.

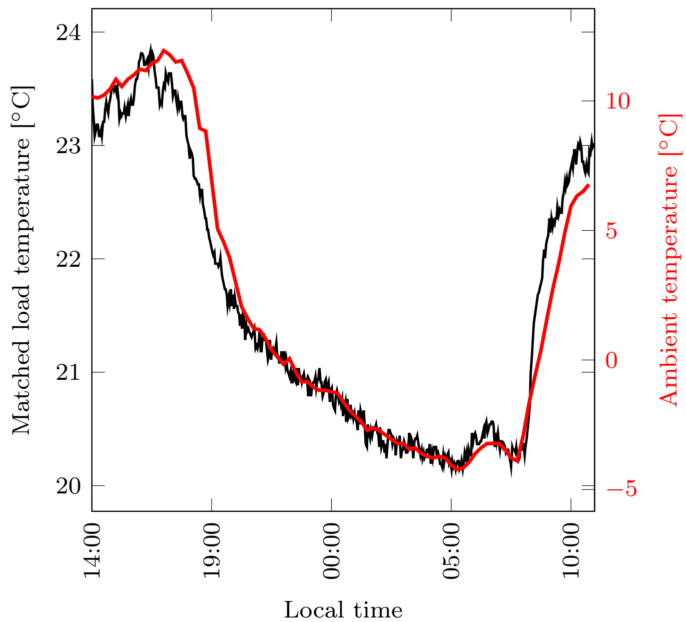

Figure 2.2: Analog signal chain temperature (black) and ambient temperature (red) during a day in mid-March, when ambient fluctuations were small relative to those in the summer months.

To assess the gain stability of the prototype, we computed the overlapping Allan deviation [6]. This metric requires a contiguous, uninterrupted time series, so we disabled the calibration cycle, replaced the antenna with a matched load, and collected 21 h of continuous data. Figure 2.3 shows the overlapping Allan deviation and the minimum calibration cadence of 1200 s required to meet the on-sky specification. The Allan deviation characterizes the fractional gain stability of the system as a function of integration time, τ. On the downward-sloping portion of the curve, the Allan deviation decreases as $\tau^{-1/2}$, consistent with white noise averaging down with integration time. Beyond the knee near $\tau \approx 750$ s, the deviation increases with τ, indicating that systematic gain fluctuations dominate on timescales longer than τ_{knee}. Because the calibration cadence was longer than τ_{knee}, we could not continuously integrate to reduce the noise variance in the monitor data. Consequently, we redesigned the monitor to meet the gain stability requirements.

In this chapter, we present the design, construction, and validation of an RFI monitor, emphasizing system gain stability for long-term RF environment characterization. In § 2.2,

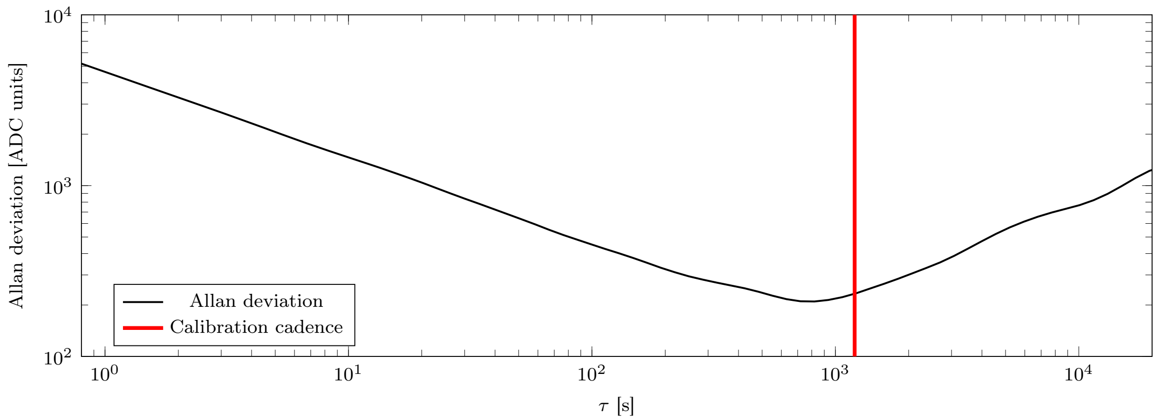

Figure 2.3: Allan deviation (black) at 1250 MHz demonstrating a stability limit at $\approx$ 750 s, beyond which systematic gain fluctuations dominate. The prototype calibration cadence of 1200 s (red) exceeded this useful integration window.

we present the requirements for the RFI monitor and detail their derivation; in § 2.3, we introduce the final design of the DRAO RFI monitor; and in § 2.4, we validate the system against these requirements. Finally, Appendix 2.A details the decisions and justifications that drove the monitor's evolution from a prototype to a commissioned instrument.

The contributions in this chapter are:

1. a rigorous requirements derivation for a wideband RFI monitoring system,
2. a novel formulation for the effect of time-varying gain drift on integrated power measurements,
3. a description of the commissioned DRAO RFI monitor design, and
4. validation measurements demonstrating that the system meets its requirements.

2.2 System Requirements

2.2.1 Frequency Coverage

The primary use case for the DRAO RFI monitor is to provide continuous coverage from 250 MHz up to 2 GHz, with defined performance targets from 350 MHz to 1800 MHz. This frequency range covers most of the operating instruments on site, as shown in fig. 2.4. These instruments include the 26 m John A. Galt telescope (JAG; 1290–1740 MHz), CHIME (400–800 MHz), the Dish Verification Antenna (DVA; 350–1050 MHz), the Canadian

Hydrogen Observatory and Radio-transient Detector (CHORD; 300–1500 MHz), the Advanced Radio Telescope Timing Array (ARTTA; 400–800 MHz, 900–1800 MHz), the Solar Flux Monitor (SFM; 2770–2830 MHz), and the Canadian Galactic Emission Mapper (CGEM; 8–10 GHz). Although continuous monitoring is a more costly and computationally demanding solution than frequency scanning, we assert that it is essential for characterizing both the continuous and transient local RF environment at the desired level of detail, consistency, and sensitivity.

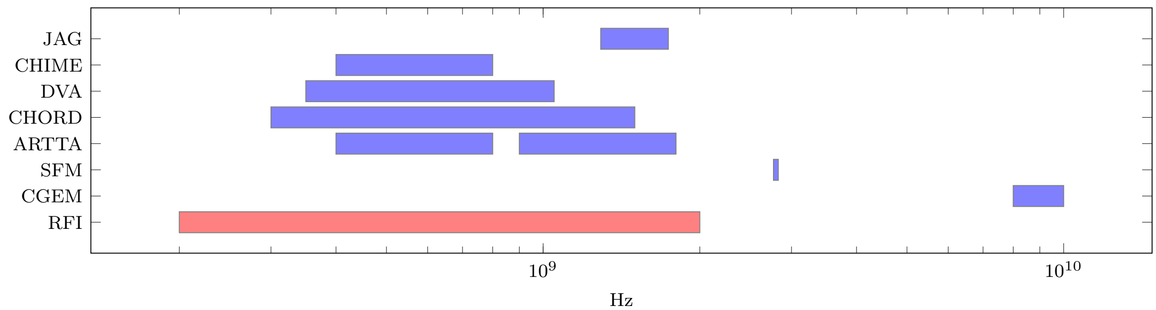

Figure 2.4: Frequency ranges of DRAO telescopes along with the RFI monitor frequency coverage. The primary frequency range of the RFI monitor covers the majority of current and proposed instruments.

2.2.2 System Temperature

Unlike a highly directional radio telescope, the monitor observes a broad environment consisting of ≈60% ground and surrounding hills and ≈40% sky, a ratio derived from the site's average horizon angle of $12.5°$ (see Appendix 2.B). According to International Telecommunication Union Radiocommunication Sector (ITU-R) Recommendation P.372-17, the cosmic background noise temperature is 2.7 K and the nominal surface temperature is 288 K [7]. Weighting these temperatures by the monitor's spatial coverage yields an expected antenna noise temperature of $T_{\text{ant}} \approx 174\,\text{K}$. Adopting a slightly conservative value of $T_{\text{ant}} = 170\,\text{K}$, we specified a matching target noise figure for the analog receiver of 2 dB ($T_{\text{rec}} \leq 170\,\text{K}$). Consequently, the target system temperature for the instrument is $T_{\text{sys}} = T_{\text{ant}} + T_{\text{rec}} = 340\,\text{K}$.

2.2.3 Antenna Position

The analog receiver chain must possess sufficient dynamic range to detect low-level interference during quiet periods while avoiding saturation from intense, highly localized RFI generated by on-site human activity. Because staff, contractors, couriers, and visitors inevitably bring RFI-generating devices on site, the RFI monitor must be positioned such that a cellular device—typically the highest-power emitter at the DRAO—transmitting from any publicly accessible location cannot saturate the receiver.

To derive the required physical standoff distance, we establish the following parameters and constraints:

1. The most powerful localized transmitters are likely to be cellular uplinks operating in the $824–849$ MHz band (centered at 836.5 MHz) with a maximum radiated power of 1 W ($P_{\text{tx}} = 30$ dBm).

2. To strictly limit nonlinearities, the RF signal chain requires at least a 10 dB safety margin below the input power 1 dB compression point (IP1dB).

3. Prior to any programmable attenuation, the analog front end incorporates two cascaded, non-programmable low-noise amplifiers (LNAs), each providing 22 dB of gain with a 1.25 dB noise figure at 1 GHz.

The primary design constraint is preventing saturation of the second-stage LNA. At 836.5 MHz, this amplifier has an output power 1 dB compression point (OP1dB) of 19 dBm and an IP1dB of -2 dBm. To maintain the 10 dB safety margin, the input power to this second stage must not exceed -12 dBm. Accounting for the 22 dB gain of the first-stage LNA, the maximum tolerable signal power arriving at the input of the receiver is $P_{\text{target}} = -2\,\text{dBm} - 10\,\text{dB} - 22\,\text{dB} = -34\,\text{dBm}$.

Consequently, the physical placement of the monitor antenna must guarantee a minimum path loss of $L_{\text{target}} = P_{\text{tx}} - P_{\text{target}} = 64\,\text{dB}$.

2.2.4 On-Sky Time

ITU-R Recommendation RA.1513-2 states that a radio astronomy observation is considered degraded when RFI causes a 5% data loss [8]. Because the monitor is specifically designed to observe RFI, we invert this metric: data loss occurs whenever the instrument is not actively recording the RF environment. To ensure a comprehensive long-term analysis, we therefore require the monitor to maintain a 95% “on-sky” time.

2.2.5 Gain Stability

Because of highly variable ambient temperatures at the DRAO, temperature-induced gain instability is a primary design concern. Contributors to this instability include the temperature instability of the signal chain, the noise contribution of the receiver, the duration of each calibration, and the cadence of the calibration cycle.

Up to this point, signal levels have been characterized by their total power, P (W). However, evaluating wideband sensitivity requires a transition to power spectral density (PSD), S_x ($\mathrm{W\,Hz^{-1}}$). Following the Johnson-Nyquist relationship, the total noise power in a channel of bandwidth Δf (Hz) is $P = kT_{\text{sys}}\Delta f$, where k is Boltzmann's constant and T_{sys} is the system noise temperature. The radiometer input power level is therefore the PSD, given by $S_x = P/\Delta f = kT_{\text{sys}}$.

The sensitivity of a total-power radiometer employing a square-law detector[1] is defined as the smallest detectable change in the input PSD level, ΔS_x ($\mathrm{W\,Hz^{-1}}$). This theoretical limit is governed by the ideal radiometer equation, which bounds the fractional sensitivity by the system's integration time and channel bandwidth:

$$\frac{\Delta S_x}{S_x} = \frac{\Delta T_{\text{thermal}}}{T_{\text{sys}}} = \frac{1}{\sqrt{\Delta f \tau}} = \frac{1}{\sqrt{N}}\,, \tag{2.1}$$

where τ is the total integration time in seconds, $N = \Delta f \tau$ is the bandwidth-time product (equivalent to the number of independent samples collected), and $\Delta T_{\text{thermal}}$ is the equivalent temperature variation in the system due to thermal noise.

To determine the maximum allowable gain change that maintains system sensitivity, we adopt the criterion from ITU-R Recommendation RA.769-2 [9]. This recommendation specifies a maximum 10% error when measuring the sensitivity limit, ΔS_x. Applying this tolerance yields a maximum allowable PSD drift of

$$\Delta S_{x,\,\text{max}} = 0.1\frac{S_x}{\sqrt{N}} = 0.1\frac{kT_{\text{sys}}}{\sqrt{N}}\,. \tag{2.2}$$

The duration over which the monitor must remain gain-stable is dictated by the calibration cadence. For the prototype, a 95% on-sky requirement (§ 2.2.4) and a $60\,\text{s}$ calibration duration (Appendix 2.A) limit the system to a maximum of three calibrations per hour. Consequently, the monitor must maintain gain stability over an interval of $1200\,\text{s}$.

To derive a specific requirement for the maximum noise variance attributable to gain

[1] Our system records complex baseband voltages and calculates the squared magnitude to yield the signal power.

drift (motivated by the RA.769-2 requirement for maximum interferer power), we adopt the assumptions outlined in [10]:

- Thermal and gain variations are independent and add quadratically:

$$\Delta T_{\text{total}}^2 = \Delta T_{\text{thermal}}^2 + \Delta T_{\text{gain}}^2 \,, \tag{2.3}$$

  where ΔT_{total} is the total noise uncertainty added to T_{sys} as a function of time and ΔT_{gain} is the uncertainty due to gain drift.

- A fractional change in the system gain produces a proportional uncertainty in the system noise temperature,

$$\Delta T_{\text{gain}} = \frac{\Delta G}{G} T_{\text{sys}} \,. \tag{2.4}$$

We therefore require that the noise uncertainty in the system attributable to gain drift, ΔT_{gain}, be no greater than 10% of the total noise uncertainty, ΔT_{total}, such that

$$\Delta T_{\text{gain}} \leq 0.1 \Delta T_{\text{total}} \,. \tag{2.5}$$

Solving eq. (2.3) for $\Delta T_{\text{thermal}}$ and substituting into eq. (2.5), we find that

$$\Delta T_{\text{gain}} \leq 0.1 \Delta T_{\text{thermal}} \,. \tag{2.6}$$

A full derivation of eq. (2.6) is provided in Appendix 2.C. We derive a simple and measurable relationship from this requirement by substituting eq. (2.4) into eq. (2.6), yielding

$$\frac{\Delta G}{G} T_{\text{sys}} \leq 0.1 \Delta T_{\text{thermal}} \,.$$

Dividing by T_{sys} gives

$$\frac{\Delta G}{G} \leq 0.1 \frac{\Delta T_{\text{thermal}}}{T_{\text{sys}}} \,.$$

Applying the radiometer eq. (2.1) yields

$$\frac{\Delta G}{G} \leq 0.1 \frac{1}{\sqrt{N}} \,.$$

Rearranging these terms provides the final requirement,

$$\frac{\Delta G}{G} \sqrt{N} \leq 0.1 \,. \tag{2.7}$$

While eq. (2.4) is standard in the literature [10, 11], it treats gain variation as a static quantity, without accounting for how gain evolves during an integration. Because thermal variance decreases as $1/\sqrt{N}$ while gain drift accumulates over time, eqs. (2.3) and (2.7) do not accurately describe the stability of an integrated measurement. To address this, we define the noise uncertainty due to gain variation as the effect of the integrated gain on the output data:

$$\Delta T_{\text{gain,integrated}} = \frac{T_{\text{sys}}}{\tau} \int_0^{\tau} \frac{\Delta G(t)}{G(0)} dt = \frac{T_{\text{sys}}}{\tau} \int_0^{\tau} \frac{G(t) - G(0)}{G(0)} dt \,, \tag{2.8}$$

where $G(0)$ is the initial system gain at time $t = 0$, $G(t)$ is the gain at time t, and τ is the integration time. By substituting eq. (2.8) into eq. (2.6) and repeating the algebraic steps used to derive eq. (2.7), we establish a more rigorous constraint for the RFI monitor:

$$-0.1 \leq \frac{\sqrt{N}}{\tau} \int_0^{\tau} \frac{G(t) - G(0)}{G(0)} dt \leq 0.1 \,. \tag{2.9}$$

2.2.6 Autonomous Operation

The monitor must autonomously collect calibrated RF measurements and populate a continuously updated database of RFI detections. The system requires no manual intervention for data collection or processing, completely eliminating the standard astronomical paradigms of observation planning and discrete measurement sets. Users query a single database interface, regardless of their specific scientific or engineering objectives. To enable the derivation of conventional metrics such as spectral occupancy [12], the backend must compute, at a minimum, a time series of integrated power spectra.

2.2.7 Software Modularity

To provide utility beyond that of a standard radiometer, the monitor is designed as an SDR. This architecture enables advanced site characterization using techniques that operate on complex-baseband waveforms or amplitude envelopes, such as automatic modulation classification, novelty detection, and RF transmitter fingerprinting. Consequently, the system transcends basic spectrum monitoring to function as a comprehensive platform for RF environment analysis.

2.2.8 Standard Interfaces

Finally, standardized connectivity is critical for an open research platform. All subsystems must interconnect using established hardware and protocols. To this end, we have specified VMEbus International Trade Association (VITA) 49.0 over User Datagram Protocol (UDP)/Internet Protocol (IP) as the radio transport protocol for time-domain streams [13]. Data transport is routed over 10 Gbit Ethernet using quad small form-factor pluggable (QSFP) connectors, which allows the system to interface directly with commercially available hardware.

2.3 System Design

2.3.1 Sensitivity

While no formal sensitivity requirement was established, an RF link budget guided component selection to maximize the sensitivity of the monitor. Although an RFI monitor cannot practically match the sensitivity of a dedicated radio telescope, benchmarking against established astronomical standards provides a necessary target for system performance.

ITU-R Recommendation RA.769-2 provides fundamental protection criteria for radio astronomy [9]. To establish the most stringent possible performance baseline, we look to the Square Kilometer Array (SKA), which will be the world's largest and most sensitive radio telescope. Because it requires protection for the deepest possible observations, the SKA RFI/Electromagnetic Compatibility (EMC) Standard [14] builds upon the [9] methodologies to derive even stricter protection levels for both continuum and spectral-line observations. Expressed as PSD ($\mathrm{dBm\,Hz^{-1}}$), these limits are

$$S_x(f_c) = \begin{cases} -17\log_{10}(f_c) - 192, & \text{continuum for } \quad 50\text{–}2000\,\mathrm{MHz}\,, \\ -17\log_{10}(f_c) - 177, & \text{spectral line for } 50\text{–}2000\,\mathrm{MHz}\,, \end{cases} \tag{2.10}$$

where f_c is the center frequency of an observation in megahertz. These protection levels assume fractional channel bandwidths of 1% of f_c for continuum observations and 0.001% of f_c for spectral-line observations.

Because the radiometer equation demonstrates that noise temperature uncertainty scales proportionally to $1/\sqrt{\Delta f}$, the sensitivity difference between these two observation modes is entirely dictated by the ratio of their bandwidths. The resulting 15 dB offset between the continuum and spectral-line protection levels is calculated directly as

$$10 \log_{10} \left(\sqrt{\frac{\Delta f_{\text{continuum}}}{\Delta f_{\text{spectral line}}}} \right) = 15\,\text{dB}\,. \tag{2.11}$$

A primary use case for this monitor is the identification of unshielded oscillators and other narrowband emitters at the DRAO. Because this task is analogous to spectral-line observing, we use the SKA spectral-line fractional bandwidth to define our system's maximum channel bandwidth. Applying the 0.001% requirement to the monitor's lowest calibrated frequency of $350\,\text{MHz}$ yields a minimum required channel bandwidth of $\frac{0.001}{100} \times 350\,\text{MHz} = 3.5\,\text{kHz}$.

In § 2.4.4, we validate the sensitivity of the commissioned RFI monitor against these targets and compare it to commercial off-the-shelf receiver systems.

2.3.2 Antenna and Structure

The requirement established in § 2.2.3 mandates a minimum path loss of $64\,\text{dB}$ between a transmitting cellular or mobile device on site and the RFI monitor. At a cellular uplink frequency of $f = 836.5\,\text{MHz}$ (wavelength $\lambda = 0.36\,\text{m}$), standard free space path loss (FSPL) requires a physical separation distance of approximately $45\,\text{m}$ to achieve this attenuation.

We chose to mount the monitor antenna on the roof of the main DRAO office building, a central location that provides a clear line of sight to most of the site's telescopes. However, because the structurally permissible installation site on this roof is less than $45\,\text{m}$ away from the parking lot and other high-traffic, publicly accessible areas on site, FSPL alone is insufficient to prevent receiver saturation. To satisfy the attenuation requirement, we utilize the metal flashing of the building's parapet as a diffracting knife-edge. With the mast's horizontal position locked by these structural limits, the antenna height becomes the sole variable parameter available to optimize the diffraction loss.

To determine the required antenna height, we construct a geometric model (fig. 2.5) using the following parameters and assumptions:

- The parapet flashing acts as a diffracting knife-edge at a height of $h_e = 8.56\,\text{m}$ above the ground.
- The height of the antenna above the parapet edge, h_a, is the variable design parameter (yielding an antenna elevation of $h_e + h_a$).

- A structural assessment dictates a fixed horizontal standoff distance from the building edge to the antenna mast of $d_{ea} = 10\,\mathrm{m}$.
- The mobile transmitter is located at a height of $h_t = 2\,\mathrm{m}$, the standard propagation model assumption for a pedestrian or vehicle.
- The horizontal distance from the transmitter to the building edge, d_{te}, can vary between 1–$100\,\mathrm{m}$.
- Ground reflections are neglected in this geometric model.

This geometry is shown in fig. 2.5. Following the conventions of [15], d_1 and d_2 represent the slant distances from the diffracting edge to the transmitter and the monitor, respectively. The angles α_1 and α_2 represent the diffraction angles between the building edge and the respective endpoints.

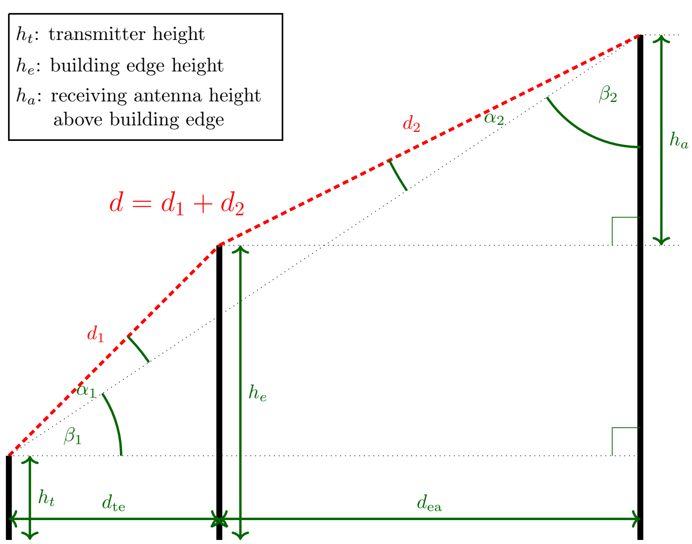

Figure 2.5: Antenna position geometric model, where h_t is the transmitter height, h_e is the height of the diffracting building edge, and h_a is the receiving antenna height above the building edge.

The total propagation loss, L, comprises two distinct components. The first is the FSPL defined by [16] as

$$20\log_{10}\left(\frac{4\pi d}{\lambda}\right),$$

where d is the total path distance. The second is the additional attenuation due to knife-edge diffraction defined by [15] as

$$6.9 + 20\log_{10}\left(\sqrt{(\nu - 0.1)^2 + 1} + \nu - 0.1\right),$$

where $\nu = \sqrt{(2d/\lambda)\alpha_1\alpha_2}$ is the dimensionless Fresnel diffraction parameter. Combining these two components yields the total propagation loss constraint:

$$L = 20\log_{10}\left(\frac{4\pi d}{\lambda}\right) + 6.9 + 20\log_{10}\left(\sqrt{(\nu - 0.1)^2 + 1} + \nu - 0.1\right). \tag{2.12}$$

Note that at the grazing angle ($\alpha_1 = \alpha_2 = 0$), the diffraction parameter reduces to $\nu = 0$ corresponding to half of the first Fresnel zone being obstructed. Under this condition, the knife-edge introduces ≈6 dB of additional attenuation. Consequently, the total loss is 6 dB greater than FSPL alone.

Because the transmitter distance (d_{te}) is highly variable, we must ensure the 64 dB attenuation requirement is met under worst-case spatial conditions. For a given antenna height (h_a), we calculate the path loss across the entire range of possible pedestrian locations ($d_{te} \in [1, 100]$ m) and identify the minimum resulting loss. Plotting the worst-case loss as a function of h_a (fig. 2.6) demonstrates that the antenna height must be constrained to $h_a \leq 4.4$ m to guarantee sufficient attenuation.

To provide complete $360°$ coverage of the terrestrial RF environment and capture all local interference events, the system requires a vertically polarized, wideband, omnidirectional antenna. We selected the passive Alaris OMNI-A0190 compact monitoring antenna, which operates from 20 MHz to 6 GHz. Across the target frequency range of 350–1800 MHz, the isotropic antenna gain is nominally 0 dBi with a minimum of -4 dBi at 625 MHz [17]. The antenna azimuth ripple is < 6 dB with an E-plane beamwidth of $60°$. The antenna's E-plane beamwidth favors the horizon, enabling monitoring of the surrounding area but limiting sensitivity to overhead emitters. Future work may address this through a temporary antenna substitution or a hemispherical replacement, allowing characterization of satellite and aircraft emissions. The installation—comprising the antenna, the mounting structure, and the delineated sections of the signal chain—is shown in fig. 2.7.

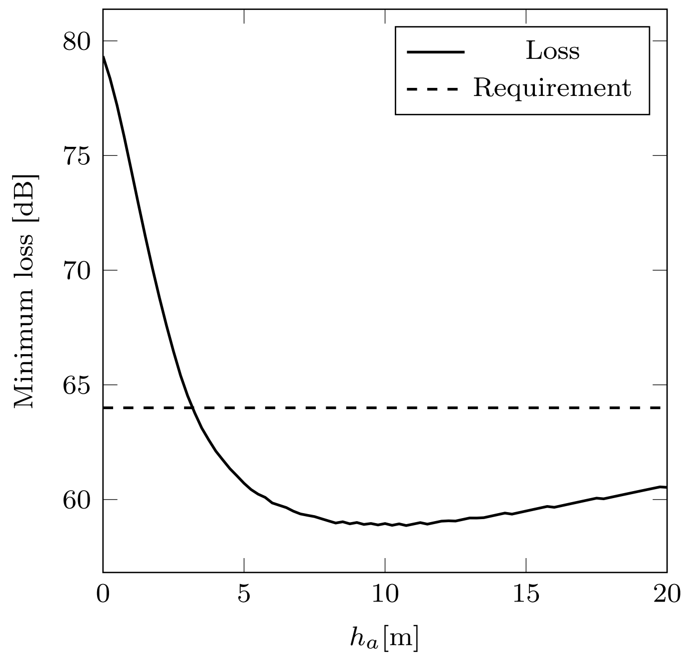

Figure 2.6: Worst-case path loss evaluated across all transmitter distances (d_{te}) as a function of antenna height (h_a). To guarantee the target loss of 64 dB, the antenna must be mounted lower than 4.4 m above the diffracting edge.

2.3.3 Mechanical Structure

To ensure temperature stability, a copper block of mass ≈ 20 kg serves as a common thermal backplane into which the RF signal chain is partially embedded. A direct-to-air thermoelectric assembly (TEA) mediates heat exchange between the copper plate and the external environment while maintaining thermal isolation.

A simplified model of the design is shown in fig. 2.8. The core design principle is to highly insulate the enclosure to approximate an adiabatic boundary, restricting thermal gradients entirely to the TEA interface. By placing the TEA control sensor between the RF components and the TEA itself, the system leverages the thermal mass of the copper. The TEA can rapidly compensate for ambient fluctuations before the thermal wave propagates through the block, ensuring the RF signal chain remains at a stable temperature. Crucially, there is no requirement for the signal chain to be held at a specific absolute temperature; the sole operational constraint is to limit the temperature-induced gain instability, which is driven by temporal temperature gradients rather than absolute thermal limits.

A 3D rendering of the signal chain enclosure is shown in Figure 2.9. The copper backplane provides a thermal standoff approximately 10.2 cm thick between the TEA interface and the embedded RF components. To enforce the adiabatic assumption, the environmen-

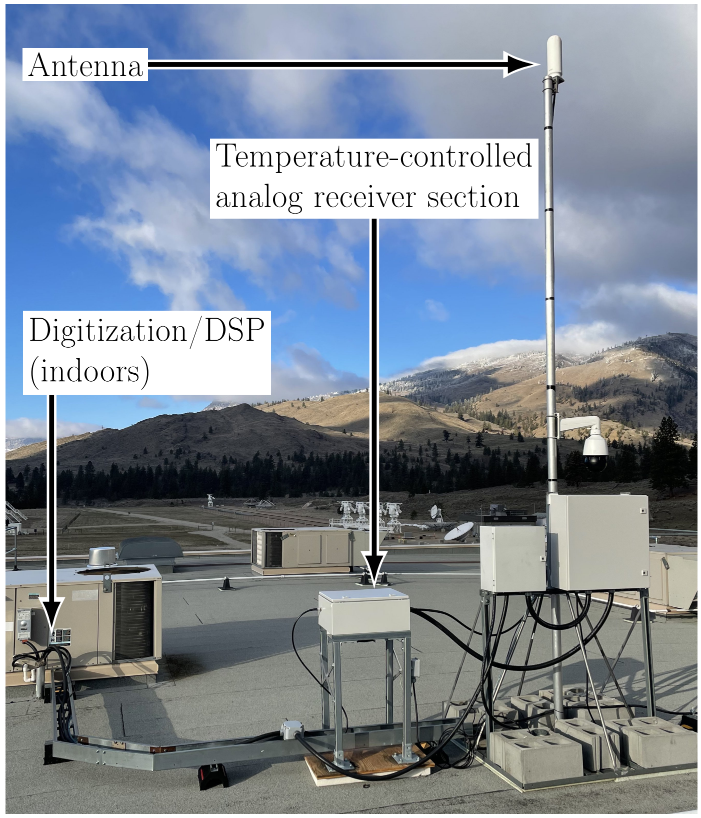

Figure 2.7: The passive Alaris OMNI-A0190 monitoring antenna mounted on the roof of the main building, alongside the delineated sections of the downstream signal chain. The mounting height is strictly constrained by the knife-edge diffraction loss required to shield the receiver from publicly accessible areas.

tally sealed steel enclosure is completely back-filled with rigid, closed-cell insulating foam (Airex T92.60). Figure 2.10 shows the signal chain and copper mass surrounded by the insulating foam. Wires are embedded into the rigid foam, and spray-foam insulation is used to prevent air gaps.

2.3.4 Calibration System

To maintain the gain stability of the analog signal chain, the monitor employs a standard Y-factor calibration. This process utilizes two discrete states to characterize the system noise: a “hot” measurement and a “cold” measurement. A simplified block diagram of this architecture is shown in fig. 2.11.

Two high-isolation switches control the signal routing for these measurements. During calibration, the first switch disconnects the antenna from the signal path and replaces

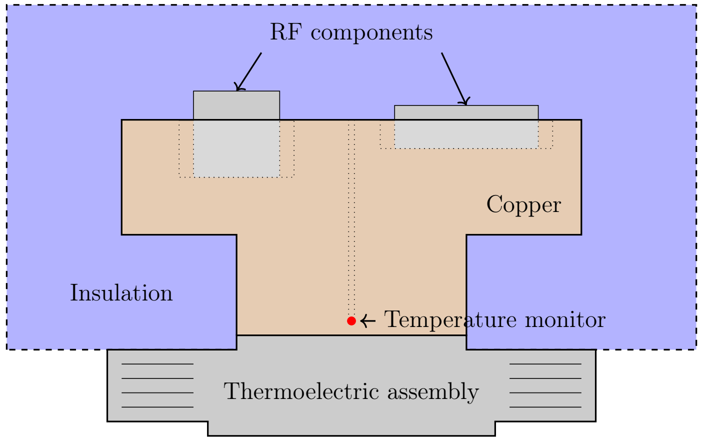

Figure 2.8: Simplified model of the insulated box containing the RF signal chain. Assuming an adiabatic boundary (dashed) and steady-state power dissipation from the RF components, heat transfer is restricted exclusively to the TEA. Because the feedback sensor is located near the TEA–copper interface, ambient temperature variations are actively rejected, and the RF components experience negligible temperature variation.

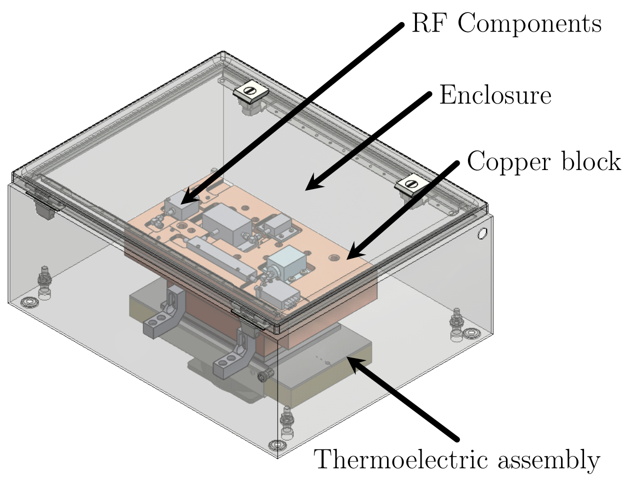

Figure 2.9: A 3D rendering of the signal chain enclosure illustrating the component layout. To minimize thermal leakage, all depicted void spaces within the physical steel enclosure are completely filled with insulating foam.

it with the primary $50\,\Omega$ matched load. A “hot” measurement is then recorded by toggling the second switch to couple the noise diode into the signal path. Conversely, a “cold” measurement is obtained by toggling the second switch to replace the noise diode with a secondary $50\,\Omega$ matched load. This dual-switch configuration ensures the receiver remains isolated from the external RF input while providing a stable reference for the Y-factor calculation. The physical layout and placement of these switching and calibration

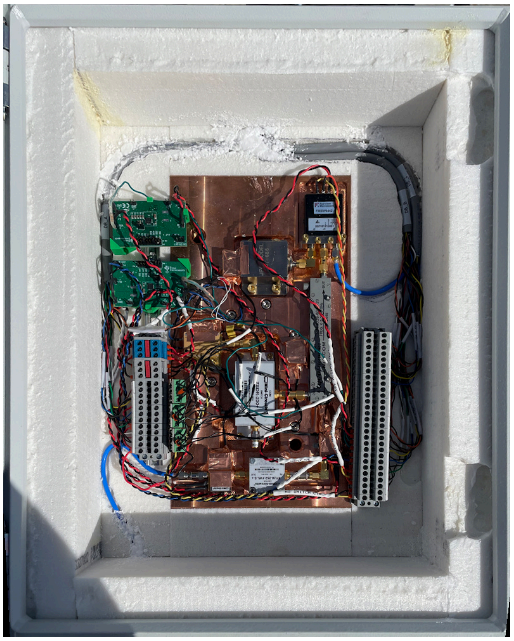

(a) The copper backplane embedded in insulating foam, with signal chain components installed.

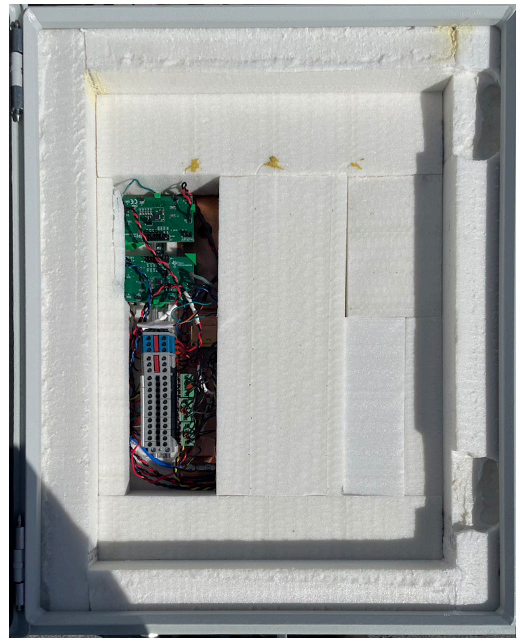

(b) The signal chain components enclosed in precision-cut insulating foam.

(c) The completed foam insert, fully enclosing the thermally stabilized signal chain.

Figure 2.10: The signal chain enclosure at successive stages of foam insulation assembly.

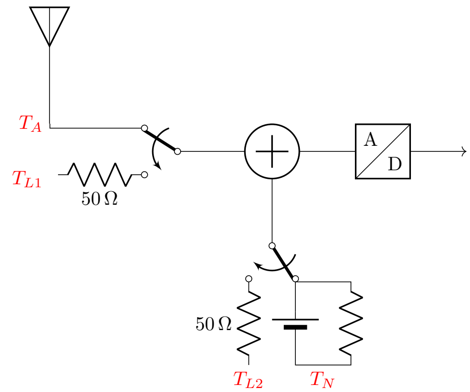

Figure 2.11: Simplified block diagram of the calibration system architecture. The first switch selects between the antenna (T_A) and a primary $50\,\Omega$ matched load (T_{L1}). The second switch toggles the calibration input between a noise diode (T_N) and a secondary $50\,\Omega$ matched reference load (T_{L2}).

components within the analog receiver are detailed in § 2.3.5.

Generating Calibrated Integrations

The "hot" and "cold" measurements yield the uncalibrated power spectra, P_{hot} and P_{cold}, respectively, in dimensionless analog-to-digital converter (ADC) units. Combined with the known physical temperature of the primary matched load, T_{L1}, and the excess noise temperature of the noise diode, T_N, these spectra allow us to derive the receiver noise temperature, T_{rec}, and the system gain, G.

As illustrated in fig. 2.11, the total equivalent noise temperatures for the two states are $T_{\text{hot}} = T_{L1} + T_N$ and $T_{\text{cold}} = T_{L1}$. We omit the contribution of T_{L2} to T_{cold}; because it is injected via a $-20\,\text{dB}$ coupler (detailed in § 2.3.5), its effective noise contribution is on the order of only $1\,\text{K}$. The ratio of the measured powers defines the standard Y-factor:

$$Y = \frac{P_{\text{hot}}}{P_{\text{cold}}} = \frac{T_{\text{hot}} + T_{\text{rec}}}{T_{\text{cold}} + T_{\text{rec}}} \,. \tag{2.13}$$

Solving eq. (2.13) for the receiver temperature yields:

$$T_{\text{rec}} = \frac{T_N}{Y - 1} - T_{L1} \,. \tag{2.14}$$

With T_{rec} established, we use the Johnson-Nyquist noise power relation with a gain term, $P = GkT\Delta f$, to determine the system gain, G. This factor maps the physical power to the digitized ADC output. Evaluating the "cold" state ($P = P_{\text{cold}}$ and $T = T_{\text{cold}}$), the gain is

$$G = \frac{P_{\text{cold}}}{k(T_{L1} + T_{\text{rec}})\Delta f}\,. \tag{2.15}$$

The units of G are ADC units per watt. Note that we could equivalently compute the gain using the "hot" state parameters (P_{hot} and T_{hot}).

This calibration gain is then applied to each uncalibrated integration, P_{uncal} (ADC units), to yield a calibrated PSD in physical units (W Hz^{-1}):

$$S_{x,\text{cal}} = \frac{P_{\text{uncal}}}{G\Delta f}\,. \tag{2.16}$$

The factor $1/G$ reverses the instrumental gain, converting the raw digitized ADC units back into absolute power. Finally, $S_{x,\text{cal}}$ is converted to logarithmic units (dBm Hz^{-1}) to allow for direct comparison against the RFI sensitivity requirements established in § 2.3.1.

Meeting the On-Sky Requirement

To maximize the system's sensitivity to transient RFI, the monitor must maintain an on-sky observing efficiency of at least 95%. Satisfying this requirement necessitates a careful balance: the calibration cycle must be as brief as possible, yet long enough to achieve an acceptable signal-to-noise ratio in the Y-factor measurement.

To determine the minimum number of averaged integrations, N, required for calibration, we assess the uncertainty in the Y-factor. Assuming the noise power in each frequency channel follows a normal distribution, the standard deviation of an averaged measurement decreases by a factor of $1/\sqrt{N}$. Let σ_{hot} and σ_{cold} represent the standard deviations of a *single* integration for the hot and cold states, respectively. If we average N_{hot} and N_{cold} integrations to form the mean power measurements P_{hot} and P_{cold}, respectively, the standard errors of these means are $\sigma_{\text{hot}}/\sqrt{N_{\text{hot}}}$ and $\sigma_{\text{cold}}/\sqrt{N_{\text{cold}}}$.

Applying the standard formula for propagation of uncorrelated errors to the Y-factor quotient ($Y = P_{\text{hot}}/P_{\text{cold}}$), we obtain the variance:

$$\left(\frac{\sigma_Y}{Y}\right)^2 = \left(\frac{\sigma_{\text{hot}}}{P_{\text{hot}}\sqrt{N_{\text{hot}}}}\right)^2 + \left(\frac{\sigma_{\text{cold}}}{P_{\text{cold}}\sqrt{N_{\text{cold}}}}\right)^2, \tag{2.17}$$

which simplifies to the final form used to evaluate the calibration duration:

$$\frac{\sigma_Y}{Y} = \sqrt{\frac{\sigma^2_{\text{hot}}}{N_{\text{hot}}P^2_{\text{hot}}} + \frac{\sigma^2_{\text{cold}}}{N_{\text{cold}}P^2_{\text{cold}}}}\,. \tag{2.18}$$

To empirically establish these parameters, we recorded 500 independent integrations of both the hot and cold calibration states (as defined in § 2.3.4). Using the sample means and standard deviations across all frequency channels, we calculated the fractional Y-factor uncertainty for symmetric integration counts ($N = N_{\text{hot}} = N_{\text{cold}}$).

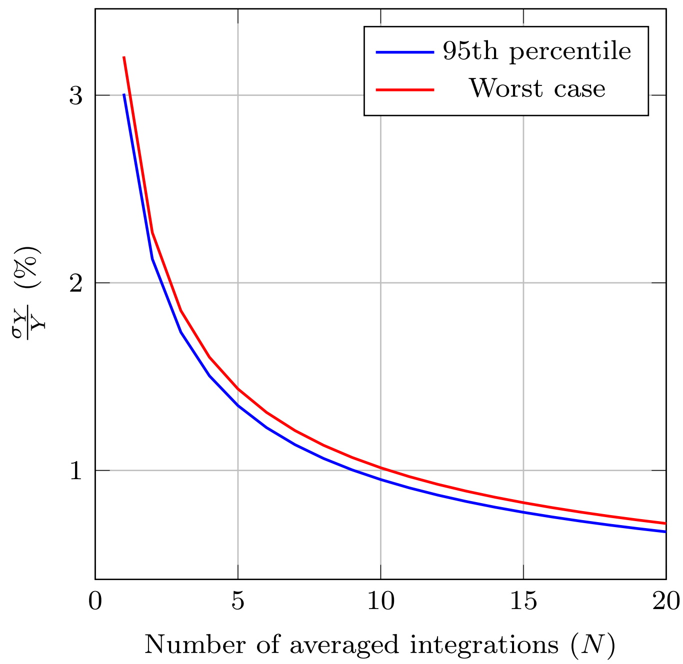

Figure 2.12: Fractional Y-factor uncertainty as a function of the number of averaged integrations (N). The curves represent the maximum (worst-case; red) and 95th percentile (blue) values across the measured spectrum. At $N = 9$, 95% of the frequency bins achieve a fractional 1σ uncertainty of less than 1%.

As shown in fig. 2.12, evaluating σ_Y/Y across all frequencies demonstrates that averaging $N = 9$ integrations ensures the 1σ uncertainty remains below 1% for 95% of the frequency bins. With an integration time of 0.75 s, recording 9 integrations for both the hot and cold states results in a total calibration duration of 13.5 s.

To ensure this 13.5 s calibration overhead consumes no more than 5% of the monitor's operational time, the system must execute a maximum of 320 calibration cycles per day. This establishes a strict calibration cadence (the duration between the start of successive calibrations) of exactly 270 s. Within this cadence, 13.5 s are spent calibrating, leaving 256.5 s for active RF environment monitoring.

To guarantee this precise 270 s cadence without being subject to the unpredictable latency of software-level execution, the analog signal chain switches are controlled directly by the ADC's general purpose input/output (GPIO) pins. The FPGA implements a hardware-level sequencer; pairs of timestamps and control values are pushed from the software into a first in, first out (FIFO) block. As the internal system clock progresses, entries are dequeued from the FIFO and applied directly to the GPIO pins, ensuring microsecond-level switching alignment with the digitized samples. The exact state of the switches is permanently embedded within the VITA 49.0 metadata trailer of each integration, allowing any actuation delays to be compensated downstream and guaranteeing seamless application of the calibration coefficients.

2.3.5 Analog Receiver

A simplified block diagram of the analog receiver is shown in fig. 2.13. The signal chain is divided into two distinct physical environments: a temperature-controlled front end deployed outdoors adjacent to the antenna mast (fig. 2.7), and an ambient-temperature section collocated indoors with the DSP subsystem.

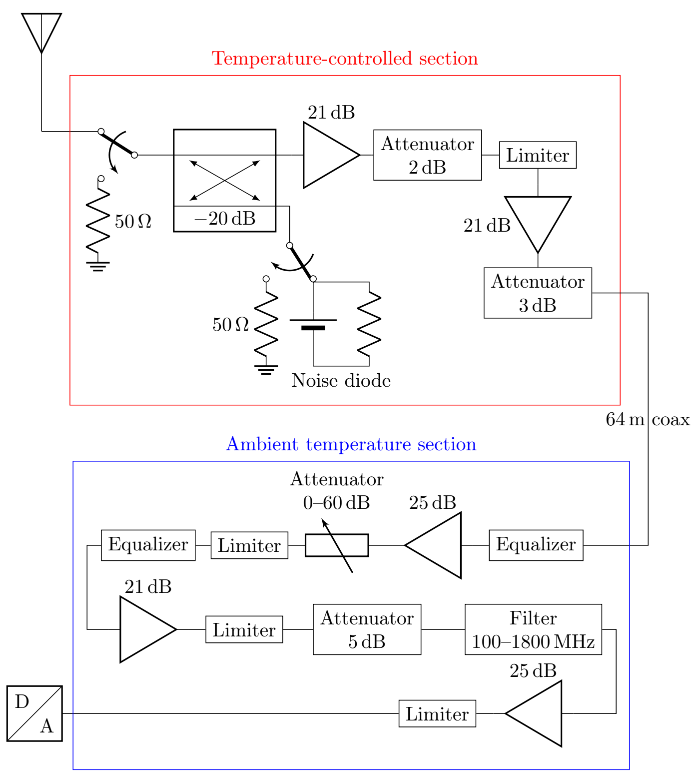

Figure 2.13: Analog receiver block diagram. The temperature-controlled section is deployed outdoors adjacent to the antenna mast and contains the calibration hardware. The ambient section is deployed indoors and contains an equalizer, adjustable gain, and an anti-aliasing filter. The sections are connected by 64 m of coaxial cable.

The temperature-controlled section contains the calibration hardware and the first-stage LNAs to minimize front-end insertion loss and ensure optimal gain stability. The primary $50\,\Omega$ matched load (JFW 50T-686) is bolted directly into the copper backplane (see § 2.3.3) to leverage its immense thermal mass. A Fairview FMSW6441 latching switch routes the signal path between the antenna and this matched load, utilizing two TTL inputs to direct the mechanical latch to either RF port. Because the switch requires no continuous holding current, it dissipates negligible heat, preserving the system's temperature stability during calibration. Furthermore, its low insertion loss of 0.3 dB minimizes its contribution to the system noise temperature.

Mechanical latching switches intrinsically emit electrical transients during actuation. To characterize this effect, high-resolution time-domain samples were captured from each RF port immediately after actuation using a 20 GSPS bandwidth oscilloscope (fig. 2.14). Testing revealed that the J1 port exhibits a significantly higher-energy transient; consequently, the quieter J2 port was selected for the antenna path. Crucially, measurements confirmed that the transient is exclusively visible on the active port, indicating no measurable leakage to the disconnected path. However, any transient emitted from the active port will couple directly into the antenna and broadcast as localized RFI. While a solid-state switch (SKY13373-460LF) was evaluated and found to produce a smaller transient, it exhibited higher insertion loss and inferior port isolation compared to the mechanical switch. The mechanical switch was ultimately selected because its transient duration is exceedingly brief (< 1 ns) and occurs at a predictable, regular cadence, making it straightforward to identify and flag in the local telescope data. The secondary reflection pulses observed on the oscilloscope were caused by a coaxial mismatch in the experimental setup. These subsequent pulses are not present in the installed system.

The signal travels indoors via 64 m of coaxial cable to the ambient-temperature section. This stage equalizes the line loss, applies programmable attenuation, boosts the signal for the ADC, and provides low-pass anti-aliasing filtering (100–1800 MHz). To manage the system gain dynamically, the control software interfaces directly with networked programmable attenuators, allowing attenuation levels to be adjusted remotely during periods of elevated RFI (such as during facility weekend tours).

To protect the amplifiers and the ADC from damage or saturation during severe RFI events, three signal limiters are distributed throughout the signal chain. The first two (Mini-Circuits ZFLM-252-1WL-S+) are used to protect the subsequent amplifiers; they feature a 0.1 dB compression point at -10 dBm and a maximum output power of 0 dBm. The final limiter (Mini-Circuits VLM-33S+) protects the ADC. It has a maximum output power of 12 dBm, and at an input power of 1 dBm, its spurious-free dynamic range (SFDR)

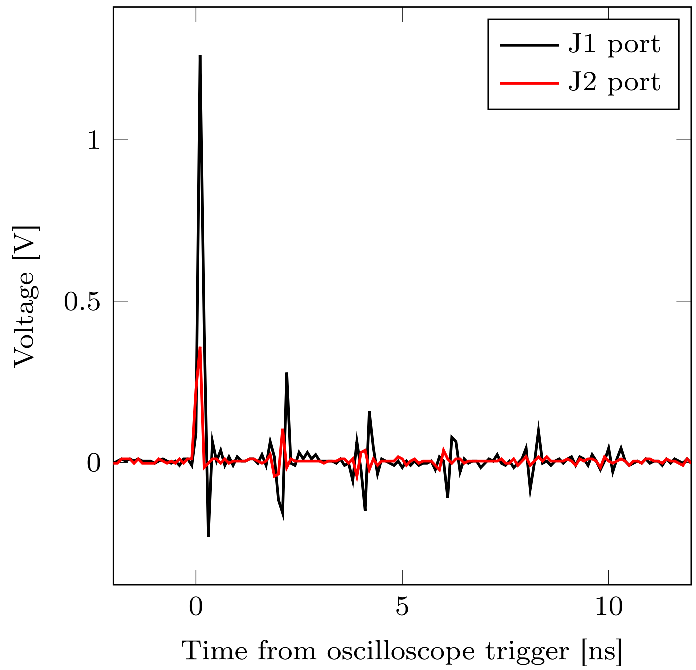

Figure 2.14: Time-domain transients generated by the latching switch at the antenna output. The J1 port exhibits a more energetic response; consequently, J2 is utilized for the antenna. The subsequent reflection pulses are an artifact of a coaxial mismatch in the experimental setup and do not occur in the installed system.

is approximately 60 dBc. For comparison, the ADC saturates at 0.96 dBm, possesses a maximum absolute input limit of 17 dBm, and has an SFDR of approximately 50 dBc. The limiters provide sufficient dynamic range and linearity without introducing spurious intermodulation artifacts into the data.

2.3.6 Digital Signal Processing

The DSP subsystem utilizes an architecture typical of a single-dish backend and largely leverages prior design work developed for other telescopes at the DRAO. This includes precise timing controls for the calibration noise source in the analog receiver, ensuring strict synchronization with the integration time of the spectrometer. Figure 2.15 illustrates the data flow through the DSP subsystem, which interfaces directly with the output of the analog signal chain (fig. 2.13). The RF signal is digitized using an interleaved ADC, such that a 4 GSPS sampling rate yields the full 2 GHz of usable bandwidth.

The first DSP stage is an FPGA-based oversampled polyphase filterbank (OSPFB) (oversampling factor $O = 16/15$) that produces 16 complex (I/Q) single-polarization channels

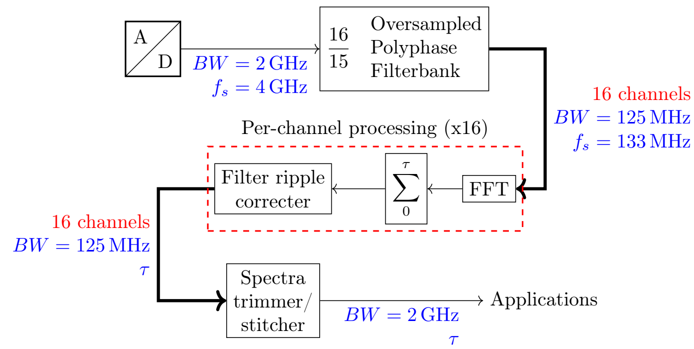

Figure 2.15: DSP block diagram detailing the 16-channel oversampled filterbank implemented in the FPGA, which feeds into 16 critically sampled channelizers implemented in the GPU.

from the digitized input. The sample rate for each channel is 133 MHz, providing a usable Nyquist bandwidth of 125 MHz. This architecture renders each of the 16 channels comparable to the bandwidth of a high-end SDR, such as the Ettus Universal Software Radio Peripheral (USRP) X310. Although the channel center frequencies are fixed, the filterbank provides contiguous, full-frequency coverage over the entire range of interest.

The OSPFB is implemented on the ICE platform [18] and relies on its native board management and control software. The internal architecture of these FPGA DSP stages is shown in fig. 2.16. The design directly reuses code from the CHIME project, as well as legacy firmware from previous instruments based on the Kermode FPGA platform [19, 20]. The ICE platform streams data to downstream compute resources via eight 10 Gbit Ethernet links. Each link carries two single-polarization channels, packetized as 16 bit+16 bit complex samples and framed according to the VITA 49.0 radio transport standard [13]. Each UDP packet contains two single-polarization channels. Crucially, all bit growth is preserved from the OSPFB through to the second-stage graphics processing unit (GPU) processing, employing a similar approach to that of [21].

The second DSP stage is a GPU-based, critically sampled fine channelizer hosted on a high-performance server. The server ingests the raw time-domain packets from the two Ethernet links using Data Plane Development Kit (DPDK) and routes the data directly to a bank of two NVIDIA GeForce RTX 2080 Ti GPUs where channelization and accumulation occur. Following accumulation, the pipeline mathematically corrects the filter ripple introduced by the first-stage OSPFB.

For each integration cycle, the resulting data products are published via 16 Zero Message Queue (ZeroMQ) streams, which are available to any subscribed service on the local

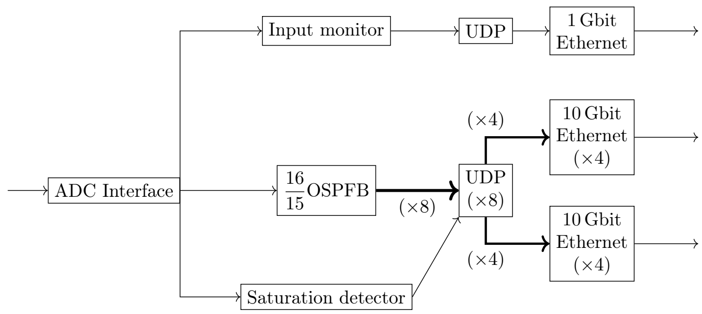

Figure 2.16: FPGA DSP block diagram showing the digitized data being processed by the OSPFB, a saturation detector, and an input monitor, all of which broadcast their outputs over UDP channels.

network. By design, the spectrometer's integration time and frequency channel bandwidth are interdependent parameters. For example, achieving a channel bandwidth of $3.33\,\mathrm{kHz}$ dictates a minimum integration time of $\approx 50\,\mathrm{ms}$ (constrained by memory bandwidth), whereas an integration time of 1 s allows for a minimum channel bandwidth of $\approx 100\,\mathrm{Hz}$ (constrained by GPU memory).

For standard operations, a channel bandwidth of $3.33\,\mathrm{kHz}$ was selected to meet the specific RFI requirements defined in § 2.2. Concurrently, an integration time of 0.75 s was chosen to manage data storage volumes. Note that these internal hardware integration times are distinct from the τ-duration integrations discussed in § 2.2.5. Because the raw spectra are output at a fast $0.75\,\mathrm{s}$ cadence, downstream applications can subsequently average these data to achieve an effective integration time of τ, or any longer duration demanded by specific science goals.

These 16 data streams can be individually subscribed to by various software applications. The primary downstream application subscribes to all 16 streams, trims the overlapping oversampled channel edges, concatenates the data, and republishes a single, unified ZeroMQ stream containing contiguous $3.33\,\mathrm{kHz}$ channels from 0–2 GHz with unified metadata. This main stream provides the definitive data product used by the majority of the facility's monitoring services.

2.3.7 Observing Software

The "publishing pipeline" continues in the observing software when a calibrated integration publisher subscribes to the concatenated, uncalibrated spectra, applies a Y-factor calibration, and republishes each integration. This architecture allows downstream software to subscribe to either calibrated or uncalibrated spectra and receive a full-bandwidth data product for every integration.

The fundamental data products generated by the system are one calibrated and one uncalibrated file every hour, each containing exactly one hour of data. These files are saved in Signal Metadata Format (SigMF), which has been modified for use with power spectra. These files are stored locally for two weeks, during which time they are available for download and can be explored using a web-based tool. This tool renders a selected hour of data as a time-frequency "waterfall" plot that can be interactively panned and scaled at full resolution. After two weeks, a reduced statistical dataset is automatically archived, and the full-resolution data files are deleted.

A primary obstacle to the success of previous RFI monitoring projects at the DRAO has been the sheer volume of data produced. Massive datasets are costly to store and often yield limited immediate insight into the local RF environment. To address this, we employ the data collection and analysis approach previously presented in [22], which represents a large number of events in a compact, easily mineable format. Using this method, individual signals are autonomously detected, and key features—such as duration, bandwidth, power, and time of day—are stored in place of the raw spectra. Future integration of the analysis technique presented in [4] will enable baseband signal reconstruction. Subsequent analysis of these baseband signal representations, using methods such as the unsupervised modulation recognition introduced in [5], will further expand the feature set for each detection.

2.3.8 Instrument Software

All of the system software is containerized, and the monitoring and control of the instrument are handled by a Tango Controls system [23]. Most telemetry data consist of temperatures collected by the LabJack T4, which are sent to a time-series database. This database is connected to a web front-end that allows users to easily visualize the telemetry data and configure temperature alerts.

A dedicated microservice monitors 23 specific "bands of interest." These encompass protected radio astronomy bands focused on single frequencies (e.g., 408 MHz and 1420 MHz), active local cellular bands (e.g., 824–829 MHz), and broad operational bandwidths for on-site telescopes (such as 400–800 MHz for CHIME). The service subscribes to the calibrated integrations and, depending on the band, computes a subset of metrics for each integration, including the mean power, maximum power, spectral occupancy, flux, and centroid. These data are also sent to the time-series database and have provided valuable insights into local cellular activity. Figure 2.17 shows a diurnal cycle of occupancy for one such band of interest. This variation is likely attributable to the routine influx of delivery drivers, employees, and contractors to the site during standard working hours.

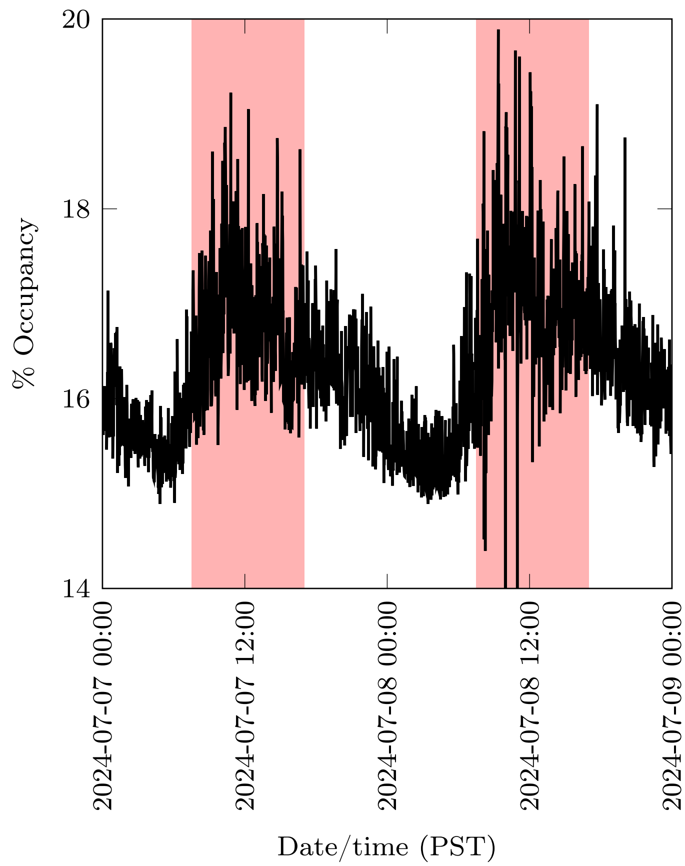

Figure 2.17: Two days of occupancy data for the $350–1050\,\mathrm{MHz}$ band, demonstrating a strong diurnal cycle. Standard working hours (07:30–17:30) are shaded in red.

2.4 Validation

2.4.1 Receiver Noise Temperature

To determine the receiver noise temperature, T_{rec}, we measured the analog signal chain using a noise-figure analyzer (NFA). Figure 2.18 demonstrates that the requirement of $T_{\mathrm{rec}} < 170\,\mathrm{K}$, established in § 2.2.2, is successfully met across the entire bandwidth. An approximately $5\,\mathrm{m}$ length of LMR-600 coaxial cable connects the antenna output to the analog signal chain input, contributing $\approx 40\,\mathrm{K}$ to the overall system noise temperature. To validate the physical measurement, we computed a comprehensive RF link budget, which predicts $T_{\text{rec}} = 134.7\,\text{K}$ at $1\,\text{GHz}$. This theoretical prediction aligns exceptionally well with the empirical measurement of $T_{\text{rec}} = 133.5\,\text{K}$ at $1\,\text{GHz}$ shown in fig. 2.18.

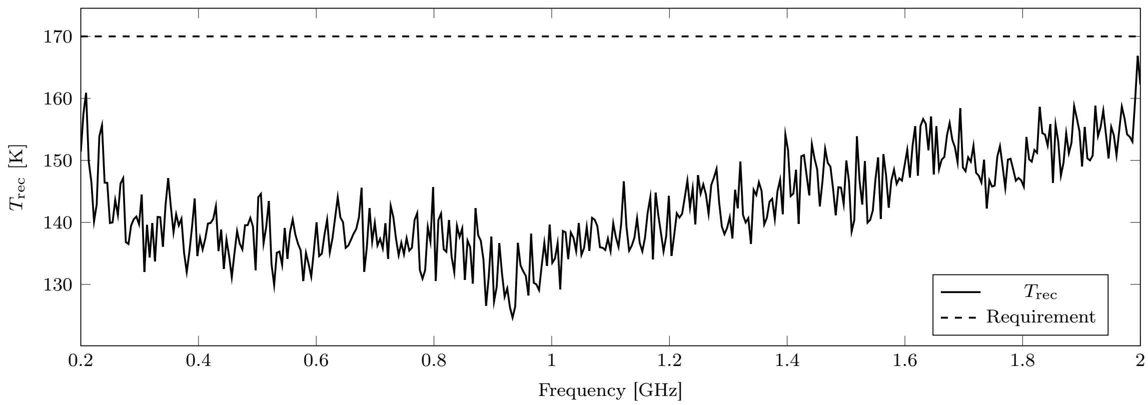

Figure 2.18: Receiver noise temperature across the operational bandwidth, demonstrating a maximum of less than $170\,\text{K}$.

2.4.2 Temperature-Induced Gain Stability

To characterize the temperature-induced gain instability described in § 2.2.5, we used an environmental chamber to sweep the ambient temperature while a vector network analyzer (VNA) recorded the corresponding gain variations.

For the $1200\,\text{s}$ comparison period defined in eq. (2.9), the most extreme temperature change recorded at the DRAO in 2024 was $3.3\,°\text{C}$ (from 26.0–$29.3\,°\text{C}$). Our objective was to verify that the system meets the stability requirement under these extreme conditions, which correspond to a constant temperature ramp rate of $0.165\,°\text{C}\,\text{min}^{-1}$. We cycled the environmental chamber across this range at this precise rate while the VNA continuously recorded the system's gain. Figure 2.19 compares these empirical measurements against the theoretical upper bound established in eq. (2.9).

2.4.3 Integrated System Stability

The final validation test was a repeated characterization of the system's Allan deviation, following the methodology described in § 2.2.5. The commissioned system was reinstalled at the same physical location on the DRAO main building roof where the prototype had been evaluated. Data were collected over a $119\,\text{h}$ period, with the input signal coming from the internal matched load rather than the antenna. The noise diode was switched in

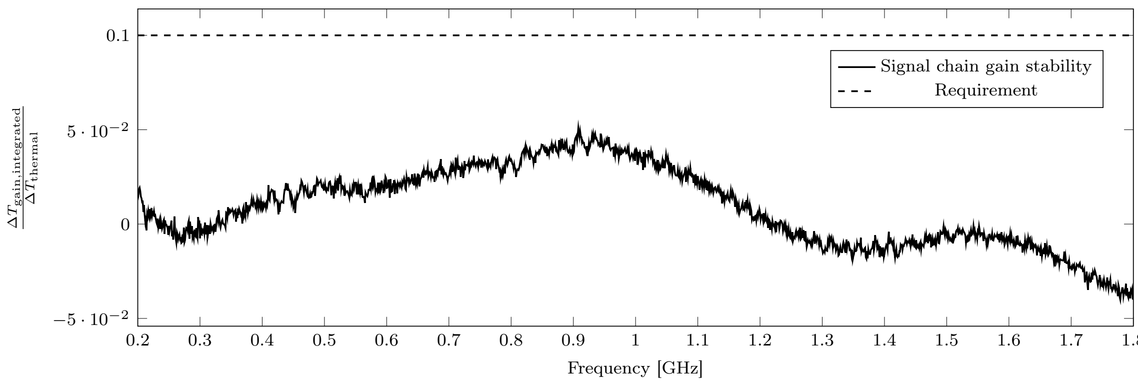

Figure 2.19: Gain stability of the system, demonstrating that the performance requirement is safely met over a $1200\,\mathrm{s}$ duration across the entire operational bandwidth.

for $6.75\,\mathrm{s}$ every $300\,\mathrm{s}$. This dataset was utilized to generate two distinct Allan deviation curves: an uncalibrated time series to determine the necessary calibration cadence, and a calibrated time series to demonstrate that long-term integration successfully reduces the noise variance.

Figure 2.20 shows the Allan deviation of the uncalibrated system. These results indicate that the system remains stable for approximately twice as long as the prototype. It is now possible to calibrate on the original $1200\,\mathrm{s}$ cadence while continuously reducing the noise variance in the measurement.

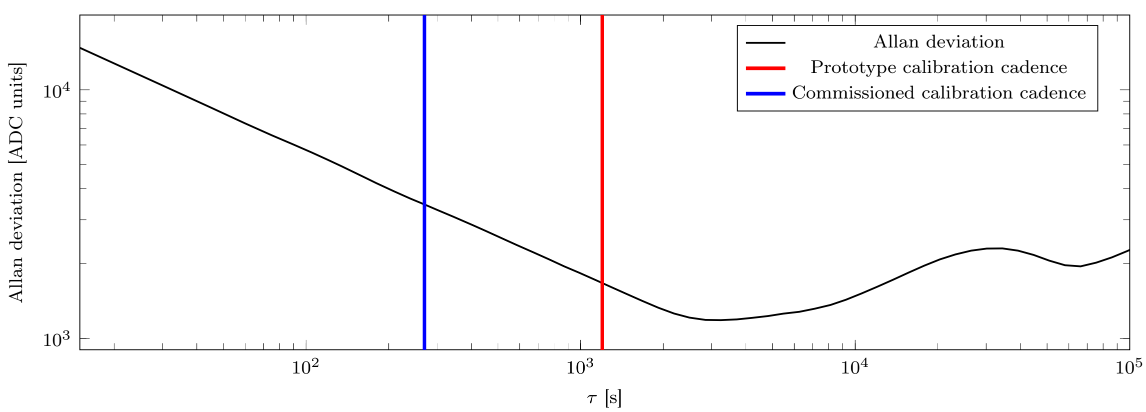

Figure 2.20: Uncalibrated Allan deviation (black) at $1250\,\mathrm{MHz}$. The RFI monitor data cannot be integrated for longer than $\approx 1400\,\mathrm{s}$ before system drift causes the noise variance to increase. The original prototype calibration cadence of $1200\,\mathrm{s}$ (red) occurs prior to this stability knee, and the minimum possible calibration cadence for the commissioned instrument ($270\,\mathrm{s}$, blue) is well to the left of the drift-limited regime.

The Allan deviation of the calibrated system is shown in fig. 2.21. These data demonstrate that the signal can be continuously integrated for at least $100\,000\,\mathrm{s}$ ($\approx 28\,\mathrm{h}$) while the noise variance continues to decline. Because no local minimum is yet apparent in fig. 2.21, no definitive limit for the maximum integration duration can be established. However, we consider a guaranteed 24 h integration period to be more than sufficient for the RF site characterization tasks envisioned in § 2.1.

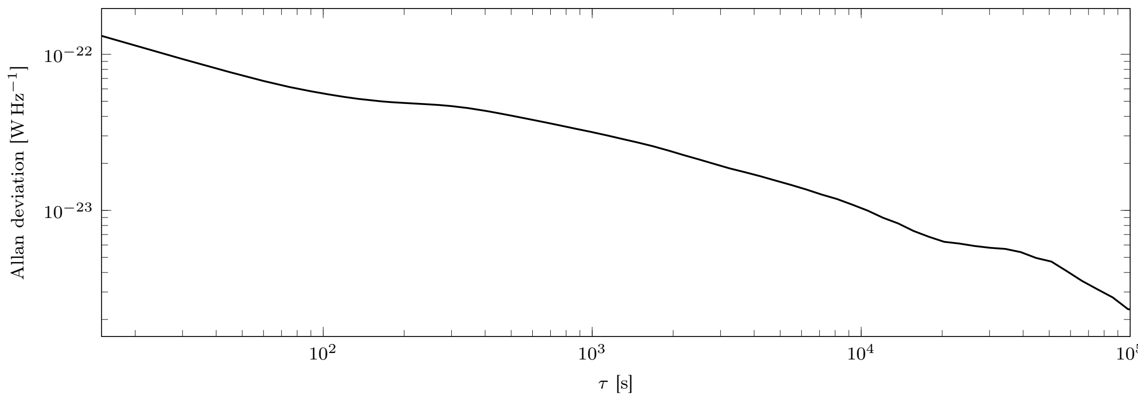

Figure 2.21: Calibrated Allan deviation at 1250 MHz, demonstrating that the RFI monitor data can be integrated up to $\approx$ 100 000 s without system drift dominating the noise variance. The slope change at $\tau = 300$ s is likely explained by a timing misalignment where the calibration cycle switching was happening in the middle of integrations, while the source of the slope at $\tau = 20\,000$ s is unknown – although possibly just a lack of data points.

2.4.4 Sensitivity

As established in § 2.3.4, the calibrated data produced by the RFI monitor are expressed as PSDs. As shown in fig. 2.22, the observed noise floor of the instrument is ≈ -172 dBm Hz^{-1}. This sensitivity is several decibels better than high-end commercial real-time spectrum analyzers, such as the Tektronix RSA7100B or Rohde and Schwarz FSWX [24, 25], demonstrating the advantage of a signal chain specifically optimized for the DRAO environment.

To validate this sensitivity, we calculated a refined estimate of the total system noise temperature, T_{sys}. Combining the measured receiver temperature ($T_{\text{rec}} = 133.5$ K) from § 2.4.1, the predicted antenna temperature ($T_{\text{ant}} = 174$ K) from § 2.2.2, and the contribution from the coaxial cable and adapters between the antenna and receiver ($T_{\text{coax}} = 84.8$ K), we find a system temperature of $T_{\text{sys}} = T_{\text{ant}} + T_{\text{coax}} + T_{\text{rec}} = 392.3$ K. Applying the Johnson-Nyquist relation established in § 2.2.5, this temperature corresponds to a predicted noise PSD of -172.7 dBm Hz^{-1}. This prediction is in excellent agreement with the observed baseline S_x shown in fig. 2.22.

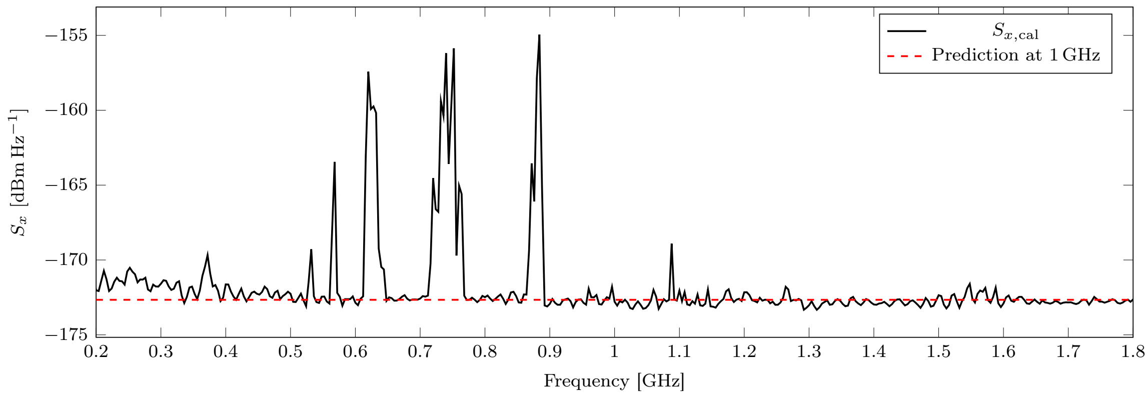

Figure 2.22: Sample spectra output from the RFI monitor showing a measured system sensitivity of ≈ -172 dBm Hz^{-1}, aligning with the theoretical prediction (red dashed line) of $S_x = -172.7$ dBm Hz^{-1}.

2.5 Conclusions and Future Work

In this work, we have presented the design, implementation, and commissioning of the new wideband RFI monitoring system at the DRAO. The instrument now provides continuous, calibrated spectral coverage from 350–1800 MHz with a system sensitivity of approximately -172 dBm Hz^{-1} and stability sufficient for integrations exceeding 24 hours; a comprehensive summary of instrument specifications is provided in table 2.1. Through rigorous design and validation, the monitor achieves the precise temperature control, strict gain stability, low system noise temperature, and rapid calibration cadence required to support long-term statistical characterization of the radio environment at DRAO.

The system significantly improves on the previous prototype in both performance and autonomy. The software-defined architecture, continuous operation, and standardized data interfaces allow it to serve as a shared infrastructure for spectrum management, instrument commissioning, and research into advanced RFI detection algorithms. The ability to correlate RFI events with site activities and environmental factors enables study of how local interference evolves over time and how it impacts scientific observations.

Planned upgrades will further expand the monitor's capabilities. A dual-input intermediate frequency (IF) system will extend coverage to higher-frequency bands of interest (2.75–2.85 GHz, 8–10 GHz, and the 2.4 GHz unlicensed band), supporting instruments such as the SFM and CGEM. An autonomous adaptive attenuation control loop will be developed to dynamically manage strong RFI events while maintaining system linearity. Moving the receiver nearer to the antenna and digitizing nearer to the receiver will reduce the amount of coax needed and could drop the system temperature by as much as

20%. On the software side, the real-time implementation of signal classification, modulation recognition, and anomaly detection will transform the monitor into a live research platform for machine-learning-based RFI mitigation.

Table 2.1: Summary of key system parameters for the commissioned DRAO RFI monitor.

Parameter	**Value / Description**
Calibrated frequency coverage	350–1800 MHz (continuous)
Instantaneous bandwidth	2 GHz
Channel bandwidth	3.33 kHz (standard; down to 100 Hz for long integrations)
Integration time	0.75 s (standard; down to 50 ms for transient characterization)
System temperature	≈ 390 K
Receiver noise temperature	≈ 133 K
Gain stability	Continuous integration up to 24 hours without drift dominating
On-sky fraction	$\geq 95\%$
Calibration cadence	270 s
Sensitivity	≈ -172 dBm Hz^{-1}

Appendices

2.A Redesign Decisions

2.A.1 Latching Switch and Matched Load

In the prototype, the antenna switch was not the latching type described in § 2.3.5. Instead, it dissipated power and generated heat when switched to the matched load. The matched load was a termination-style component that mounted directly onto the switch and had little thermal mass. As a result, the matched-load temperature would rise during the calibration, influencing the gain calibration results. Figure 2.23 shows the matched-load temperature during switching for both the old and new configurations, demonstrating that the matched-load temperature remains stable during switching because the switch is now latching and the new matched load has a much larger thermal mass tied into the copper backplane.

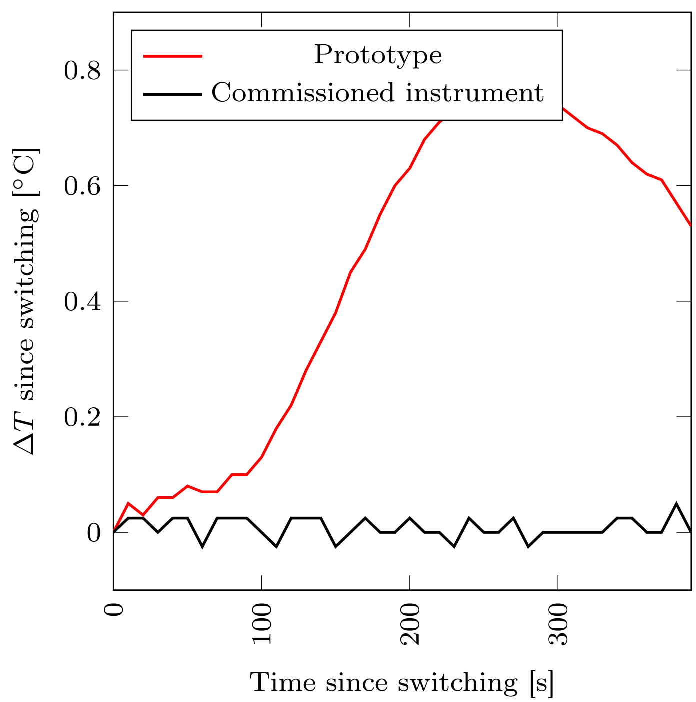

Figure 2.23: Matched-load temperature during switching before (red) and after (black) the redesign, showing that the matched load no longer heats up.

2.A.2 Directional Coupler

We also replaced the directional coupler, which injects the $20\,\mathrm{dB}$ noise-diode calibration signal into the RF chain. The original coupler had a low intrinsic insertion loss but occupied a large footprint and required SMA-to-N-type adapters at both ends to mate with the system's SMA cables. The replacement—a Tyco Electronics 2026-6012-20 coupler—features native SMA connectors and a smaller form factor, albeit with a higher intrinsic insertion loss. However, because the combined loss of the original coupler and its adapters exceeded that of the standalone Tyco unit, the new configuration yields a net performance gain. This reduction in overall signal path loss accounts for the slight improvement in T_{rec} shown in fig. 2.18. Furthermore, the more compact signal chain fits into a smaller volume, making it easier to thermally stabilize.

2.A.3 Mechanical Structure

In the prototype, an approximately $5\,\mathrm{mm}$-thick aluminum sheet served as the thermal backplane, and an air-to-air TEA interfaced with the external environment to control the signal-chain temperature. As fig. 2.2 shows, this system could not maintain a stable thermal environment. To resolve this stability issue, we redesigned the entire mechanical structure, incorporating a large copper block and the direct-to-air TEA described in § 2.3.3.

Furthermore, the loosely packed polystyrene foam insulation used in the original system was replaced with precision-cut Airex T92.60. This substantially reduces the air gap around the internal components.

We tested the thermal stability of both the original and improved mechanical designs in an environmental chamber. Figure 2.24 shows the matched-load temperature for both systems as the ambient temperature varied. In both tests, the TEA had a temperature setpoint of $24\,^{\circ}\mathrm{C}$.

2.A.4 Receiver Noise Temperature

The updated signal chain yields an $\approx 10\,\mathrm{K}$ improvement in receiver noise temperature over the prototype, as shown in fig. 2.25. We attribute this reduction in T_{rec} to the upgraded analog receiver components—specifically the new latching switch, the replacement directional coupler, and the elimination of the coaxial adapters.

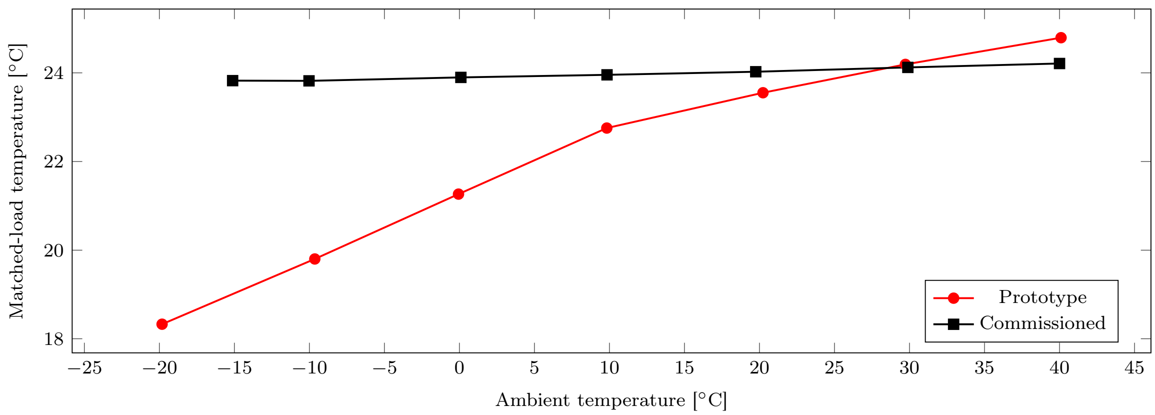

Figure 2.24: Matched-load temperature versus ambient air temperature. The mechanical redesign improves the thermal stability of the RF signal chain by more than a factor of two.

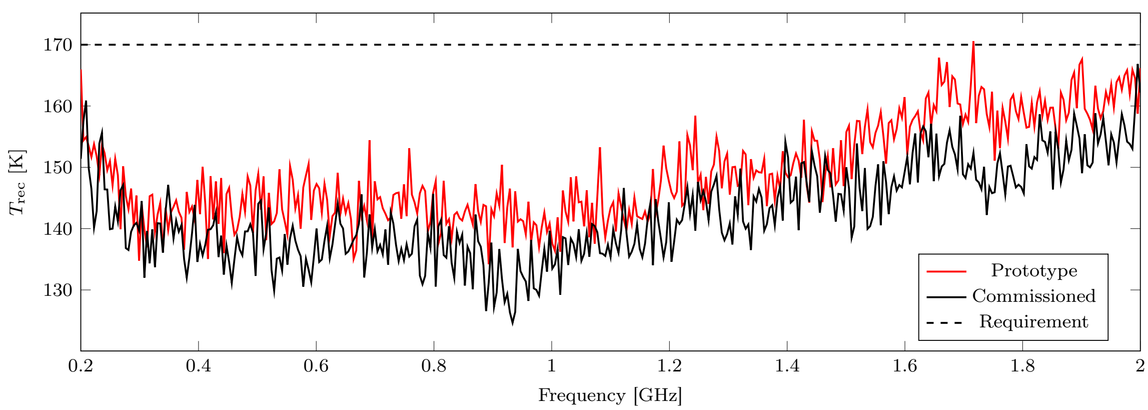

Figure 2.25: Receiver noise temperature for both the prototype (red) and the commissioned monitor (black) showing an $\approx 10\,\mathrm{K}$ improvement.

2.A.5 Temperature Signal Stability

We monitor temperatures inside the signal-chain enclosure—specifically at the matched load, the noise diode, each LNA, the copper surface, and the internal ambient air—using thermistors and the LabJack T4 digitizer. Originally, the unconditioned thermistor signals passed through D-subminiature RFI filters on the RF-shielded enclosure before reaching a resistor-divider conditioning circuit. However, testing revealed that these RFI filters and a D-subminiature connector adapter coupled significant noise into the measurements. To resolve this, we relocated the conditioning circuit inside the signal-chain enclosure, ensuring that only the higher-level, conditioned voltages pass through the RFI filters. This modification successfully eliminated the coupled noise.

2.A.6 Temperature-Induced Gain Stability

To determine if the prototype met the gain-stability requirement outlined in § 2.4.2, we evaluated its performance to establish a baseline before implementing any hardware modifications. Figure 2.26 plots the integrated gain uncertainty defined by eq. (2.9) to compare the system's performance before and after we redesigned the mechanical structure of the monitor.

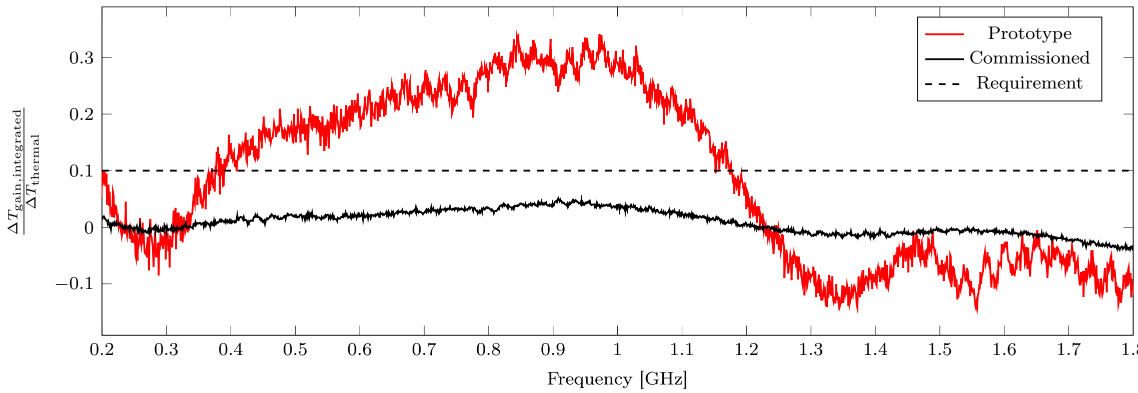

Figure 2.26: Gain stability comparison before (red) and after (black) changes were made to the system, showing that the requirement is now met for a $1200\,\mathrm{s}$ duration.

2.A.7 Calibration System

The calibration system described in § 2.3.4 significantly reduced the calibration duration compared to the prototype. The original design relied on a networked LabJack T4 to control both the antenna-to-matched-load switch and the noise-diode switch. This reliance on network commands introduced timing uncertainties during switching events. To compensate, we had to extend the duration of each calibration stage (hot and cold) and discard the data acquired immediately before and after each transition. Consequently, the overall calibration process was inherently slow.

2.B Sky-Fraction Calculation

In § 2.2.2, we state that the monitor receives noise from an environment that is $\approx 60\,\%$ ground and $\approx 40\,\%$ sky. The local horizon has an average hill elevation angle of $\alpha = 12.5°$.

Let R be the radius of a sphere, and let θ be the elevation angle from the equator. The radius of a horizontal slice at elevation θ is $R\cos\theta$, yielding a circumference of $2\pi R\cos\theta$.

A differential strip at this elevation has an area $dA = (2\pi R\cos\theta)(Rd\theta) = 2\pi R^2\cos\theta d\theta$. Integrating this from the local horizon α to the zenith yields the exposed sky area:

$$A_{\text{above}}(\alpha) = \int_{\alpha}^{\pi/2} 2\pi R^2 \cos\theta d\theta = 2\pi R^2(1 - \sin\alpha)\,. \tag{2.19}$$

Dividing the result of eq. (2.19) by the total surface area of a sphere, $A_{\text{total}} = 4\pi R^2$, provides the fractional solid angle of the sky:

$$\frac{A_{\text{above}}}{A_{\text{total}}} = \frac{1 - \sin\alpha}{2}\,. \tag{2.20}$$

Thus, for the measured average elevation angle $\alpha = 12.5°$, the fraction of the sky visible to the monitor is

$$\frac{1 - \sin(12.5°)}{2} = 39.1\,\% \approx 40\,\%\,. \tag{2.21}$$

2.C Gain-Stability Algebra

We require the uncertainty attributable to the gain drift, ΔT_{gain}, to be no greater than 10% of the total system noise uncertainty, ΔT_{total}. This condition is expressed as

$$\Delta T_{\text{gain}} \leq 0.1\Delta T_{\text{total}}\,. \tag{2.22}$$

Assuming that the thermal noise, $\Delta T_{\text{thermal}}$, and the gain-drift noise are independent, they add in quadrature to produce the total uncertainty:

$$\Delta T^2_{\text{total}} = \Delta T^2_{\text{thermal}} + \Delta T^2_{\text{gain}}\,. \tag{2.23}$$

Substituting the expression for total uncertainty into the initial requirement yields

$$\Delta T_{\text{gain}} \leq 0.1\sqrt{\Delta T^2_{\text{thermal}} + \Delta T^2_{\text{gain}}}\,. \tag{2.24}$$

Squaring both sides eliminates the radical and results in

$$\Delta T^2_{\text{gain}} \leq 0.01\left(\Delta T^2_{\text{thermal}} + \Delta T^2_{\text{gain}}\right)\,. \tag{2.25}$$

Distributing the fractional multiplier and isolating the ΔT_{gain} terms on the left-hand side of the inequality, we find

$$\Delta T_{\text{gain}}^2 - 0.01\Delta T_{\text{gain}}^2 \leq 0.01\Delta T_{\text{thermal}}^2\,, \tag{2.26}$$

which simplifies to

$$0.99\Delta T_{\text{gain}}^2 \leq 0.01\Delta T_{\text{thermal}}^2\,. \tag{2.27}$$

Dividing by 0.99 and taking the square root of both sides results in the upper bound:

$$\Delta T_{\text{gain}} \leq \frac{0.1}{\sqrt{0.99}}\Delta T_{\text{thermal}}\,. \tag{2.28}$$

Because $1/\sqrt{0.99} \approx 1.005$, setting the coefficient to 0.1 provides a conservative approximation of this upper limit accurate to within 0.5%:

$$\Delta T_{\text{gain}} \lesssim 0.1\Delta T_{\text{thermal}}\,. \tag{2.29}$$

This inequality defines the strict error budget for the instrument: as long as temperature-induced gain fluctuations are suppressed to a full order of magnitude below the radiometric white noise, the system's total measurement uncertainty will remain fundamentally bounded by integration time rather than hardware instability.

References

[1] The CHIME Collaboration, M. Amiri, K. Bandura, A. Boskovic, T. Chen, J.-F. Cliche, M. Deng, N. Denman, M. Dobbs, M. Fandino, S. Foreman, M. Halpern, D. Hanna, A. S. Hill, G. Hinshaw, C. Höfer, J. Kania, P. Klages, T. L. Landecker, J. MacEachern, K. Masui, J. Mena-Parra, N. Milutinovic, A. Mirhosseini, L. Newburgh, R. Nitsche, A. Ordog, U.-L. Pen, T. Pinsonneault-Marotte, A. Polzin, A. Reda, A. Renard, J. R. Shaw, S. R. Siegel, S. Singh, R. Smegal, I. Tretyakov, K. Van Gassen, K. Vanderlinde, H. Wang, D. V. Wiebe, J. S. Willis, and D. Wulf, "An overview of CHIME, the Canadian Hydrogen Intensity Mapping Experiment," *The Astrophysical Journal Supplement Series*, vol. 261, p. 29, 7 2022.

[2] T. Gunaratne, S. Harrison, T. Robishaw, and B. Carlson, "Characterization of the RFI environment at the DRAO: The classical approach," in *Proceedings for RFI 2022*, 2022.

[3] S. Harrison, A. Lakhani, J.-P. Lavigne, N. Bruce, T. Robishaw, R. Messing, C. Baard, M. A. Smith, B. Robert, and D. Del Rizzo, "RFI monitor system design for the Dominion Radio Astrophysical Observatory," in *2021 XXXIVth General Assembly and Scientific Symposium of the International Union of Radio Science (URSI GASS)*, (Rome, Italy), pp. 1–4, IEEE, 8 2021.

[4] N. Bruce, S. Harrison, and P. Driessen, "DDC Pool: Efficient down-conversion of signals for RFI monitoring," in *Proceedings for RFI 2022*, 2022.

[5] N. Bruce, B. Moa, S. Harrison, and P. F. Driessen, "SigCLR: A contrastive learning approach to unsupervised modulation recognition and novelty detection," in *Proceedings of Radio Frequency Interference Conference — PoS(RFI2024)*, vol. 471, p. 025, 2025.

[6] W. J. Riley, "Handbook of frequency stability analysis," Tech. Rep. NIST Special Publication (SP) 1065, National Institute of Standards and Technology, Boulder, CO, 2008.

[7] ITU-R, "Radio noise," Tech. Rep. P.372-17, ITU-R, 2024.

[8] ITU-R, "Levels of data loss to radio astronomy observations and percentage-of-time criteria resulting from degradation by interference for frequency bands allocated to the radio astronomy service on a primary basis," Tech. Rep. RA.1513-2, ITU-R, 2015.

[9] ITU-R, "Protection criteria used for radio astronomical measurements," Tech. Rep. RA.769-2, ITU-R, 2003.

[10] B. F. Burke, F. Graham-Smith, and P. N. Wilkinson, *An Introduction to Radio Astronomy*. Cambridge University Press, 4 ed., July 2019.

[11] F. T. Ulaby, R. K. Moore, and A. K. Fung, *Microwave remote sensing : active and passive*. Remote sensing ; no. 2-4, Reading, Mass: Addison-Wesley Pub. Co., Advanced Book Program/World Science Division, 1981.

[12] ITU-R, "Spectrum occupancy measurements and evaluation," Tech. Rep. SM.2256-1, ITU-R, 2016.

[13] VITA developers, "VITA radio transport (VRT)," Tech. Rep. VITA 49.0, ANSI, 2015.

[14] A. J. Otto, F. Di Vruno, T. Nkawu, P. Dewdney, G.-H. Tan, H. Smith, T. Tzioumis, and J. Jonas, "SKA RFI/EMC standard," Tech. Rep. SKA-TEL-SKO-0000202 Rev. 4, SKAO, 5 2021.

[15] ITU-R, "Propagation by diffraction," Tech. Rep. P.526-15, ITU-R, 2019.

[16] ITU-R, "Calculation of free-space attenuation," Tech. Rep. P.525-4, ITU-R, 2019.

[17] Alaris Antennas, *Compact Monitoring Antenna*, 12 2024.

[18] K. Bandura, A. Bender, J. Cliche, T. de Haan, M. Dobbs, A. Gilbert, S. Griffin, G. Hsyu, D. Ittah, J. M. Parra, *et al.*, "ICE: a scalable, low-cost FPGA-based telescope signal processing and networking system," *Journal of Astronomical Instrumentation*, vol. 5, no. 04, p. 1641005, 2016.

[19] S. T. Harrison, T. Robishaw, and G. Hovey, "A software-defined spectral line observing system for the DVA1 telescope," in *2018 2nd URSI Atlantic Radio Science Meeting (AT-RASC)*, pp. 1–4, IEEE, 2018.

[20] S. T. Harrison, G. J. Hovey, B. Veidt, Z. Ljusic, and T. Burgess, "A digital beamformer for the advanced focal array demonstrator (AFAD)," in *Millimeter, Submillimeter, and Far-Infrared Detectors and Instrumentation for Astronomy IX*, vol. 10708, p. 1070834, International Society for Optics and Photonics, 2018.

[21] G. Hobbs, R. N. Manchester, A. Dunning, A. Jameson, P. Roberts, D. George, J. A. Green, J. Tuthill, L. Toomey, J. F. Kaczmarek, S. Mader, M. Marquarding, A. Ahmed, S. W. Amy, M. Bailes, R. Beresford, N. D. R. Bhat, D. C.-J. Bock, M. Bourne, M. Bowen, M. Brothers, A. D. Cameron, E. Carretti, N. Carter, S. Castillo, R. Chekkala, W. Cheng, Y. Chung, D. A. Craig, S. Dai, J. Dawson, J. Dempsey, P. Doherty, B. Dong, P. Edwards, T. Ergesh, X. Gao, J. Han, D. Hayman, B. Indermuehle, K. Jeganathan, S. Johnston, H. Kanoniuk, M. Kesteven, M. Kramer, M. Leach, V. Mcintyre, V. Moss, S. Osłowski, C. Phillips, N. Pope, B. Preisig, D. Price, K. Reeves, L. Reilly, J. Reynolds, T. Robishaw, P. Roush, T. Ruckley, E. Sadler, J. Sarkissian, S. Severs, R. Shannon, K. Smart, M. Smith, S. Smith, C. Sobey, L. Staveley-Smith, A. Tzioumis, W. van Straten, N. Wang, L. Wen, and M. Whiting, "An ultra-wide bandwidth (704 to 4 032 MHz) receiver for the Parkes radio telescope," *PASA*, vol. 37, p. e012, Apr. 2020.

[22] S. Harrison, R. Coles, T. Robishaw, and D. Del Rizzo, "RFI novelty detection using machine learning techniques," in *RFI2019: Coexisting with Radio Frequency Interference*, 2019.

[23] Tango Controls software community, "Tango Controls." `https://tango-controls.org`, 2025. Accessed: 2025-12-30.

[24] Tektronix Inc., *RSA7100B Spectrum Analyzer Datasheet.* Tektronix Inc., 12 2025.

[25] Rohde and Schwarz GmbH & Co. KG, *FSWX Signal and Spectrum Analyzer.* Rohde and Schwarz GmbH & Co. KG, 11 2025.

Chapter 3

Visibility Domain Direction-of-Arrival Optimization for Arbitrary 2-D Arrays

> What do we call visible light? We call it color. But the electromagnetic spectrum runs to zero in one direction and infinity in the other, so really, children, mathematically, all of light is invisible.
>
> Anthony Doerr, All the Light We Cannot See

3.1 Introduction

Determining the DoA of radio signals incident on an antenna array is useful in many applications including the localization of RFI, remote sensing, spectrum management, and next-generation communications systems. The standard techniques for DoA estimation, such as beamforming, MUSIC, and Estimation of Signal Parameters via Rotational Invariance Technique (ESPRIT), emerged in parallel with the development of synthesis imaging in the RA community. As the name implies, synthesis imaging is intended to generate images of the radio sky using antenna arrays, but the formulation upon which it is built can be applied to DoA estimation as well.

These DoA methods attempt to estimate DoA by determining the linear transformation between the incident signals and those received by the antennas as represented by a steering matrix. Unlike these methods, the visibility-modelling approach casts the cross-correlations between signals received by the array elements as samples of the two-dimensional Fourier transform of the incident radiation field. Estimating the direction of arrival of a signal is then equivalent to inverting this transformation, often done by interpolating samples onto a monotonic grid and taking the 2-D fast Fourier transform (FFT), but achieved here through least-squares model fitting.

In this chapter, we show that compared to other methods, our proposed method is less computationally intensive, more noise resilient, better suited to near-horizon signals, and more accurate with fewer antennas.

The chapter is structured as follows: in section 3.1.1, we review existing DoA methods as well as radio astronomy synthesis imaging techniques; in section 3.2, we present our method; and in section 3.3, we validate the performance of our method on both simulated and measured data collected using the LWA. The incident signal is narrowband for both simulated and measured results.

3.1.1 Literature

DoA estimation methods that work for truly arbitrary array geometries are either computationally expensive, such as MUSIC [1], or suffer from low angular resolution, such as conventional and Capon beamforming [2]. Faster variants of these algorithms exist, but rely on specific array geometries as in root-MUSIC [3, 4]. Other well-known methods such as ESPRIT [5] are based on pre-defined array geometry and are still often simplified to decrease computational load [6].

Much of the literature on these methods describes proposed algorithms as computationally efficient and suitable for arbitrary geometries, then validates them exclusively on linear or rectangular arrays—often with missing elements used to justify the claim of arbitrary geometry [7]. We do not consider such arrays truly arbitrary; some authors more accurately call them *non-uniform* [8], a term we prefer. Methods dependent on specific array geometries are not reviewed further, as we are concerned only with truly arbitrary arrays.

For a method to work on an arbitrary array it should work for a random array. Random arrays are necessitated by a need to observe distant sources as in radio astronomy synthesis imaging (discussed below). Every baseline vector between two antennas that is repeated will yield a duplicate observation providing no further information about the source. Thus, effort is spent minimizing the number of repeating baseline vectors during the array design process. There are several methods used for building random arrays; one example uses a random position generator with constraints on maximum distance in any direction and minimum distance between elements [9]. Then, each randomly generated array configuration is assessed for beam size/efficiency, radiation pattern, aperture efficiency, etc. This is often balanced with the cost of laying coaxial cable to each element. An example of a random array designed using this method is shown in fig. 3.1. Another random array design method for sparse (minimum element separation of $\lambda > 1$) arrays uses Poisson Disk sampling to achieve a similar effect as the random position generator while reducing the need for trial and error [10].

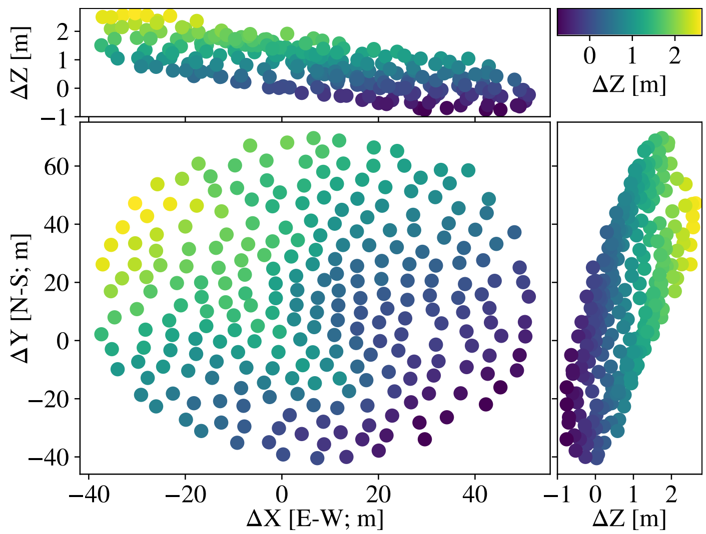

Figure 3.1: Random array geometry for the LWA-SV in New Mexico, a radio astronomy telescope.

Synthesis imaging enables the creation of all-sky images using antenna arrays such as high-resolution radio telescopes [11]. This method relies on correlating the signals between each pair of antennas to obtain samples of the so-called visibility function.

Visibility is a function of array baseline coordinates ($V(u, v)$) whereas incident radiation intensity is a function of the direction cosines $l,\ m$ as defined in eq. (3.1).

$$l = \cos(\theta)\cos(\phi) \qquad m = \cos(\theta)\sin(\phi) \tag{3.1}$$

The direction cosines are mappings from the celestial dome (described by elevation and azimuth) onto a flat plane. This relationship is illustrated in fig. 3.2.

Just as traversing between the time and frequency domains is central to much of signal processing, moving between the incident radiation intensity and visibility domains is key to synthesis imaging [11, 12]. The theory of synthesis imaging and the Fourier transform relationship between these domains is reviewed briefly in appendix 3.A. The inversion of this transform to generate images or estimate properties of the incident radiation field is achieved in all-sky imaging through convolutional resampling of visibility samples onto

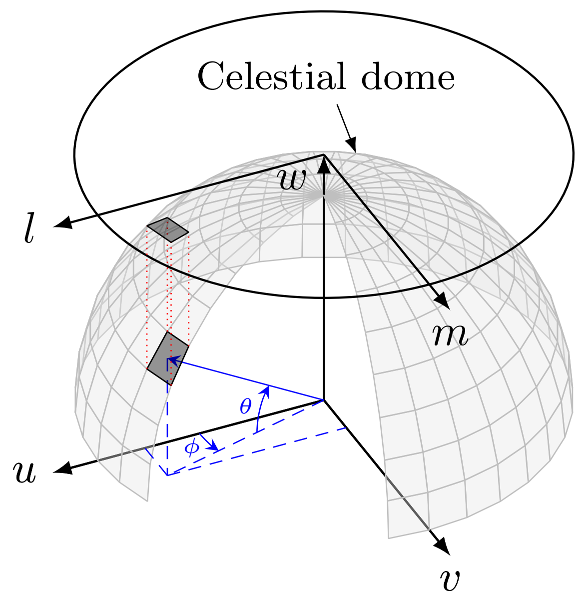

Figure 3.2: The relationship between elevation (θ), azimuth (ϕ) and (l,m) coordinates.

a regular grid followed by an FFT.

Note that while a baseline spacing of 0.5λ may be optimal for traditional beamforming, for synthesis imaging many independent baselines are needed to fill the u, v space. This is also shown in appendix 3.A. The number of baselines required changes depending on the application: wideband radio telescopes such as the LWA use hundreds of stands, yielding hundreds of thousands of unique baselines for high-angular-resolution space science, yet much smaller 4-element arrays have used synthesis imaging to study ionospheric transients and structure [13].

In order to avoid the artifacts introduced by gridding, radio astronomers have also developed advanced techniques to identify cosmological entities such as galaxies in visibility domain data [14]. These techniques are computationally slow and premised on extracting useful information about the entities such as source position, shape, and flux parameters by fitting detailed models to the data. In contrast, we simplify the extracted information to only source location and thus do not require highly detailed visibility modelling.

3.2 Methods

To determine signal DoA, a parameterized model of the signal in the visibility domain is generated and its difference from the correlator output is minimized to determine the best-fit visibility model parameters. These parameters correspond to the source's position in the sky and can then be converted into the signal's DoA.

3.2.1 Constructing a Visibility Model

We model the signal received at the array as a point source with intensity I_0 and location $(l_0,\ m_0)$. The intensity and visibility domain models are then

$$I_\delta(l, m;\, l_0, m_0) = I_0 \delta(l - l_0,\ m - m_0) \tag{3.2}$$

$$\begin{aligned} V_\delta(u, v;\, l_0, m_0) &= \iint I_\delta e^{-j2\pi(ul+vm)}\, dl\, dm \\ &= I_0 e^{-j2\pi(ul_0+vm_0)} \end{aligned} \tag{3.3}$$

3.2.2 Fitting to Sampled Visibilities

Fitting the model shown in eq. (3.3) to the visibilities output from the correlator (each correlator output is an integration, which is equivalent to a single snapshot in the language of MUSIC literature), is cast as a least-squares minimization problem. The cost function has two parameters, l_0 and m_0, and its minimum with respect to these parameters determines the best approximation of the source's location.

For each of K measured points coming out of the correlator (in which case there are K baselines in the array) the residual r_k is the difference between that measured value and the model. The least squares objective function is then the sum of the squares of the residuals at each measured visibility.

$$\begin{aligned} R(l_0, m_0) &= \sum_K \left[r_k(l_0, m_0)\right]^2 \\ &= \sum_K \left|V(u_k, v_k) - V_\delta(u_k, v_k, l_0, m_0)\right|^2 \end{aligned} \tag{3.4}$$

To minimize this non-linear function, we elect to use the Levenberg-Marquardt algorithm for its ubiquity in the literature [11]. Figure 3.3 shows a contour plot of the objective function that the Levenberg-Marquardt algorithm must traverse to find the minimum at $(l \simeq 0.45, m \simeq 0.45)$ for a signal at $(\theta = 0,\ \phi = \frac{\pi}{4})$. The function has steep gradients leading to the minimum but a near-flat domain otherwise as shown in fig. 3.3a. This requires an initial estimate (l_i, m_i) accurate to $0.4 \approx \sqrt{(l_i - l_0)^2 + (m_i - m_0)^2}$ so that the Levenberg-Marquardt has a gradient to traverse. This is not appropriate for a general direction-finding method, and we propose a solution in section 3.2.3.

3.2.3 Constructing a Domain-Wide Visibility Model

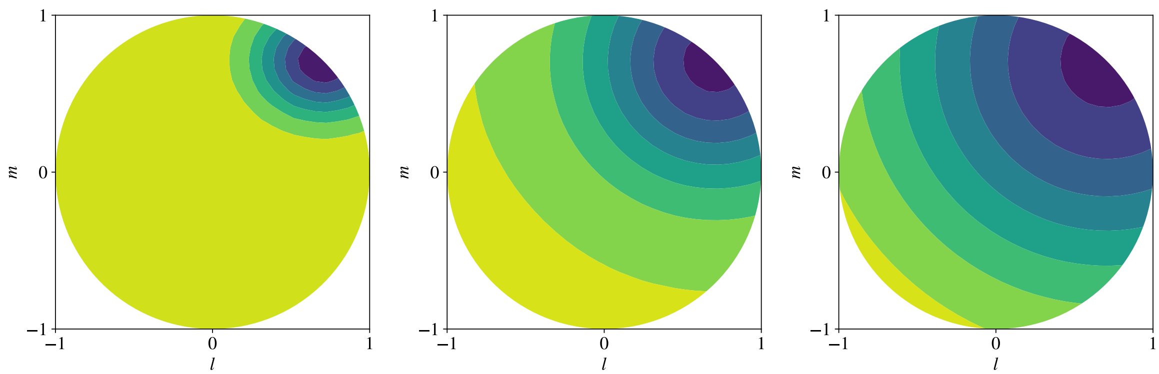

(a) Fitting signal to a point source model.

(b) Fitting signal to a Gaussian ($\sigma = 0.5$) source model.

(c) Fitting signal to a Gaussian ($\sigma = 1$) source model

Figure 3.3: Cost function domains that must be minimized to fit different models to a signal at $\theta \approx 0$ (i.e., at the horizon). The cost function for a point source model shows a flat domain except near the minimum whereas the Gaussian models show domain-wide gradients.

We bypass the need for an initial estimate by showing that the shape of the model we use during fitting does not necessarily need to be a close match to the shape of the data for the purpose of direction finding. Although our received signal is best modelled by a point source (eqs. (3.2) and (3.3)), for the purpose of DoA identification, we only need the peak of the model to align with the peak of the radiation intensity in the l, m domain. To this end we reconstruct our model so that the cost function is only flat at the minimum across the entire optimization domain.

We identify a 2-D Gaussian as a good choice for a domain-wide model with a defined peak that the Levenberg-Marquardt function can easily traverse.

$$I_{\text{Gauss.}} = I_0 \exp\left(\frac{-(l-l_0)^2-(m-m_0)^2}{2\sigma^2}\right) \tag{3.5}$$

$$\begin{aligned} V_{\text{Gauss.}} &= \iint I_{\text{Gauss.}} e^{-j2\pi(ul+vm)}\, dl\, dm \\ &= I_0 e^{-2\pi\sigma^2(u^2+v^2)+j2\pi(ul+vm)} \end{aligned} \tag{3.6}$$

The width of the Gaussian is defined by the width parameter σ such that the full width at half maximum (FWHM) is

$$\text{FWHM} = \sigma\sqrt{8\ln 2} \tag{3.7}$$

It is clear from fig. 3.3c that an initial estimate of $(0, 0)$ is always sufficient to fit a signal to a wide ($\sigma = 1$) Gaussian model since the cost function domain is never flat. Whether a point-source model is used or a wide Gaussian model, this method only works for a single source since the objective function is traversed to the minima nearest the initial estimate. There are visibility-domain deconvolution methods such as the widely used CLEAN algorithm, which effectively handle multiple sources by iteratively locating and subtracting out sources [15]. Our method could be extended to handle multiple sources by doing a similar operation set to CLEAN, namely:

1. finding the center of a source,
2. fitting 2-D Gaussian parameters to the source,
3. subtracting the source out of the visibility domain data,
4. repeating for the next located source until there are no more sources (by either knowing the number of expected sources or defining a 2-D Gaussian fit threshold for $\sigma > 1$ meaning it has fit to the residual noise in the visibilities).

Via this extension, no *a priori* knowledge of the number of sources is needed. In this sense our proposed algorithm parallels the multi-signal capability of MUSIC. While the computed MUSIC spectrum has a peak for each received signal, a threshold must be defined above which a peak is considered a signal. This is trivial if the number of received signals is known but not otherwise. Similarly, the above algorithm can be run a pre-defined number of times if the number of signals is known but if it is not a Gaussian fit width threshold can be chosen above which the next fitted peak is considered noise.

3.3 Results

To present our proposed method we first simulate data and validate the method accuracy against frequency, noise, number of antennas, integration time, and DoA. In all cases the model fitting is done using a 2-D Gaussian with $\sigma = 1$ as described in section 3.2.3. The methods chosen for comparison are MUSIC and all-sky imaging. While root-MUSIC and other MUSIC variants are computationally less expensive, we elect to use the original MUSIC algorithm for comparison as it has no array-geometry dependency and none of the variants provide greater accuracy than it.

We note that the literature surrounding MUSIC uses the term "snapshot" where we use the term integration (as does the radio astronomy field). While MUSIC calculates DoA based on a number of snapshots, we always calculate it based on a single integration. The controlled parameter in our method is the integration time.

Each simulation has different parameters but in all cases the signal is an unmodulated carrier wave as described by eq. (3.2) or eq. (3.6) depending on whether the simulation takes place in the intensity domain or the visibility domain.

3.3.1 Simulation

For the following validations an unmodulated carrier wave with a planar wavefront was simulated and its phase at each antenna was computed from the antenna positions. These time series were run through the correlator to obtain the visibilities. To do so an array geometry is required and we use the antenna positions for the Long Wavelength Array Sevilleta Station (LWA-SV) available via the LWA Software Library (LSL) [16]. The benefit to using the LWA-SV array geometry is that we can directly compare our simulated results with measured data collected using the array (section 3.3.2). The LWA-SV random array geometry is shown in fig. 3.1.

Accuracy is measured as the angle between the input vector and the predicted/calculated DoA vector. This can either be the $\cos^{-1}$ of the dot product between the two vectors, or equivalently the magnitude difference between their l, m plane projections.

Table 3.1 lists the simulations carried out and the values of the varied parameter for each. Each simulation is then discussed in detail below.

Varied parameter	Varied parameter range
Signal center frequency	4–88 MHz
Signal frequency offset	0–100 kHz
Integration time	30–0.00512 s
Antenna spacing	0.5–2.6 λ
Signal elevation	0–90°
Noise power	-40–0 dB input SNR

Table 3.1: Simulation parameters for validating the proposed DoA method.

Accuracy versus Frequency/Offset

LWA-SV is built for operation from 4–88 MHz and in a noiseless simulation our method has zero error caused by frequency across this bandwidth.

Our method relies on *a priori* knowledge of the received signal frequency and in a similarly noiseless simulation there is zero error for an offset between signal frequency and model fitting frequency of up to 100 kHz. Since the LWA-SV has a bandwidth of 100 kHz[1] this means our method does not see any error attributable to frequency offset using this array geometry.

Accuracy versus integration time

LWA-SV returns a frame of samples from each antenna every 0.00512 s. Any integration time greater than this that is a multiple of 0.00512 s is therefore easy to simulate using the LSL. When simulated in a noiseless environment there is zero error in accuracy using integration times ranging from $30 - -0.00512$ s.

Accuracy Versus Number of Antennas

A primary reason we use the LWA to validate our method is the high antenna count covering a large area ($\approx 100\,\mathrm{m} \times 110\,\mathrm{m}$). While the LWA consists of 520 dual-polarized antennas, we only make use of the 247 single polarization antennas oriented in the north-south direction. We only use a single polarization because our method depends on the spatial distribution of antennas and including both polarizations recorded at each point in space does not change the number of baselines we have to work with. Further we are interested in the DoA and not the polarization properties of the signal.

[1] The LWA bandwidth is 100 kHz in Transient Buffer Narrow (TBN) mode. The telescope has wider bandwidths (other modes) available that were not necessary for our work.

To simulate different array configurations we drop antennas from the correlator input and process those remaining. Figure 3.4 shows the accuracy curve as the overall array radius is held constant but the minimum baseline is increased. To do this ten different arrays were generated to satisfy each minimum baseline spacing requirement and the result averaged. The signal was set to the phase center of the array ($\theta = 90°$). This is shown for a 10 MHz signal; at this frequency, MUSIC fails at $\approx 1.35\lambda$ (≈ 6 antennas) while all-sky imaging and our method continue to work until the minimum of 3 antennas (the constant error term for all-sky imaging is caused by the grid cell located around the array phase center). This result is consistent with the very good array performance obtained with only 4 elements in [13]. Not shown are similar results at 5 MHz and 15 MHz where all-sky imaging and our method continue to work down to the minimum of 3 antennas while MUSIC fails at 5 antennas. Note that the average number of antennas is not an integer. Multiple array configurations (with differing performance characteristics and number of antennas) are possible at each minimum spacing and we mitigate these effects by averaging 10 arrays at each minimum spacing. There is extensive literature discussing synthesis array optimization in the context of RA [11].

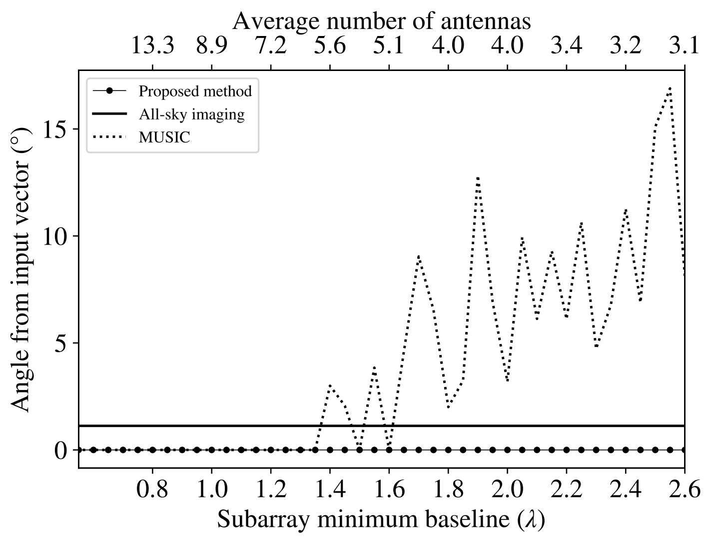

Figure 3.4: Accuracy versus average number of antennas and corresponding minimum baseline length.

Accuracy Versus Signal Elevation

To compare the accuracy of this method as the signal DoA moves around the sky we step a simulated source from the array phase center ($\theta = 90°$) to the horizon ($\theta = 0°$) and assess the accuracy at each step. For both MUSIC and all-sky imaging, as the signal approaches the horizon the sides of the celestial dome project onto a progressively smaller area of the $l,\ m$ plane (fig. 3.2), causing both an increasing distance between data points (grid cell centers) and a DoA bias towards the phase center as shown in fig. 3.5. This bias is not present with the proposed method. The saw-toothed profile of the all-sky imaging and MUSIC curves are due to the quantized l, m plane.

This results in the proposed method providing a substantial improvement over both MUSIC and all-sky imaging for low-elevation angle (θ near 0) signals. Signals of terrestrial origin refracted by the ionosphere will often arrive at low elevation angles. Thus, an accurate DoA estimate in this range of angles is important for ionospheric studies.

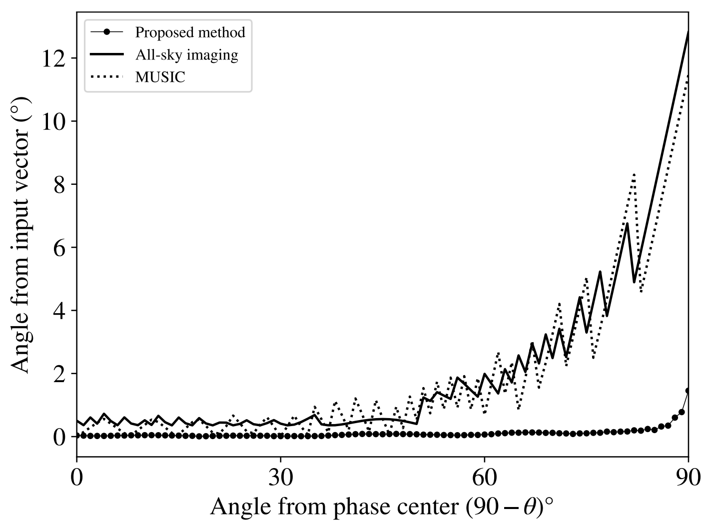

Figure 3.5: Accuracy versus signal elevation showing very little degradation in DoA accuracy near the horizon for our method while all-sky imaging and MUSIC show substantial error near the horizon.

Accuracy Versus Noise

In order to add noise to the simulation, white noise was added to the time series at each antenna as described by eq. (3.8) where a is the time series peak amplitude for the complex signal $s(t) = ae^{j2\pi ft}$, snr is the desired signal-to-noise ratio (SNR) as a linear ratio, SNR$_{dB}$ is the desired SNR in decibel and the resulting σ_I and σ_Q are the noise powers added to the real and imaginary components of the time series respectively.

$$\sigma_I = \sigma_Q = \sqrt{\frac{2a^2}{\text{snr}}} = \sqrt{\frac{2a^2}{10^{\frac{\text{SNR}_{dB}}{10}}}} \tag{3.8}$$

This allows us to show accuracy versus SNR per sample at the correlator input as shown in fig. 3.6, a valuable metric for those operating in a communications system context.

With a 0.05 s correlator integration time we see that the accuracy against noise relation only becomes relevant at signal SNRs less than -20 dB. At signal powers below -20 dB all-sky imaging outperform our method which in turn outperforms MUSIC. However for signals with SNR greater than -23 dB our method is more accurate. Our method is fractionally worse than MUSIC for SNRs greater than -28 dB. This simulation was done at the phase center of the array ($\theta = 90$) where it is apparent that all-sky imaging has a residual error, another artifact of the quantized l, m plane.

Below -36 dB (at 11° error) our method stops returning DoA results and starts returning an out-of-bounds error. This means the model fit to a location where either of l or m is greater than 1. This begins to occur for all-sky imaging below -39 dB and does not occur in this SNR range for MUSIC. We believe this is a benefit of our method versus MUSIC and all-sky imaging as low accuracy results are clearly indicated.

3.3.2 Measured Data

A 30 W radio transmitter was used to generate continuous waveform (CW) signals with known DoAs received at the LWA-SV. It was operated from Santa Fe, New Mexico under the experimental radio license WK2XTU. The bearing across the Earth's surface between the transmitter and LWA-SV gives an azimuth based on geometry of 27.65°. In practice, the azimuth may deviate from the geometric value as it reflects off the ionosphere. Due to these deviations, we do not have a ground truth for the measured data, however the simulations performed in section 3.3.1 give us confidence that the method is accurate. In addition, it is reasonable to assume the signal should have an azimuth near the geometric bearing.

MUSIC was omitted while processing hours of measured data because it proved prohibitively slow. Both our method and all-sky imaging make use of the very efficient LSL correlator, unlike our relatively naive MUSIC implementation. Furthermore, MUSIC requires calculation of a pseudo-spectrum over the grid, an operation that is computation-

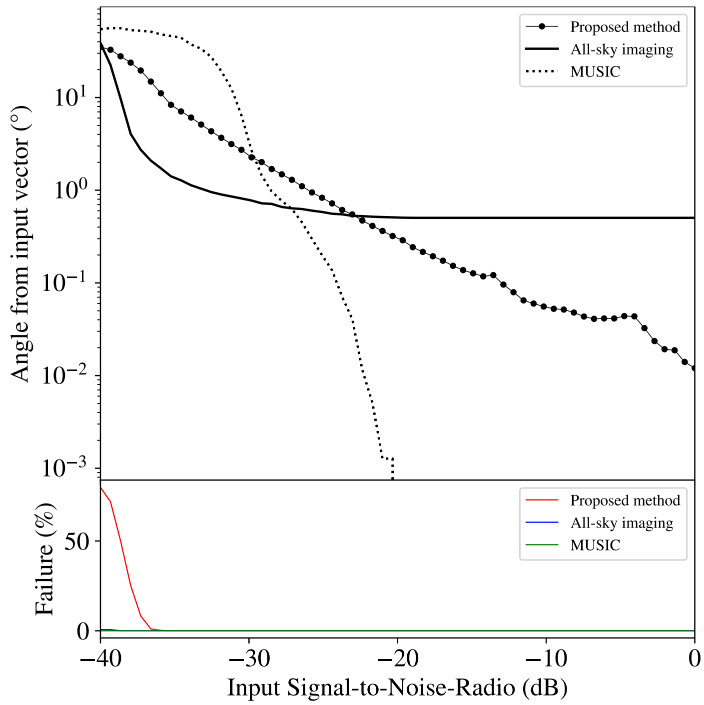

Figure 3.6: The upper plot shows accuracy versus SNR averaged over 100 runs with an 0.05 s integration time. The accuracy of the proposed method becomes worse than that of all-sky imaging near -22.5 dB SNR. The lower plot shows out-of-bound errors indicating that our method prefers returning a failure rather than a low accuracy result.

ally intensive.

Figure 3.7 shows the result of all-sky imaging and the proposed method for several hours of data recorded at LWA-SV. Results show close agreement with the geometric azimuth. The difference between mean calculated azimuth and geometric azimuth ranged from $0.20 - -1.63°$ for all-sky imaging and from $0.19 - -1.10°$ for our proposed method. The average standard deviation for all-sky imaging and our method closely agree in both azimuth and elevation as shown in table 3.2. The deviations from geometric azimuth can be attributed to ionospheric perturbations and channel effects.

	All-sky imaging	Proposed method
σ_{az}	$0.081°$	$0.082°$
σ_{el}	$0.048°$	$0.055°$

Table 3.2: Standard deviation around the calculated mean of azimuth and elevation for both all-sky imaging and the proposed method.

The gaps in the data appear where the received signal has an SNR low enough to cause an out-of-bound error (see section 3.3.1) or during periods where the experimental call sign WX2XTU is transmitted in Morse code for identification purposes.

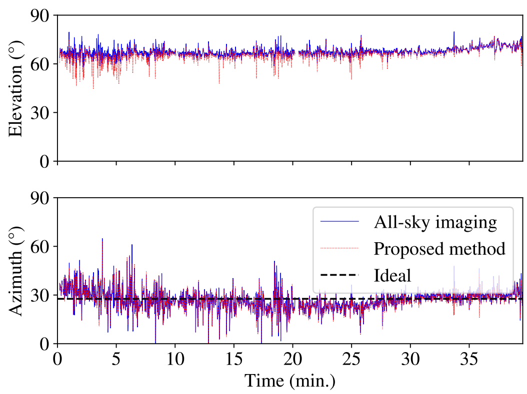

(a) 3.9995 MHz transmission frequency.

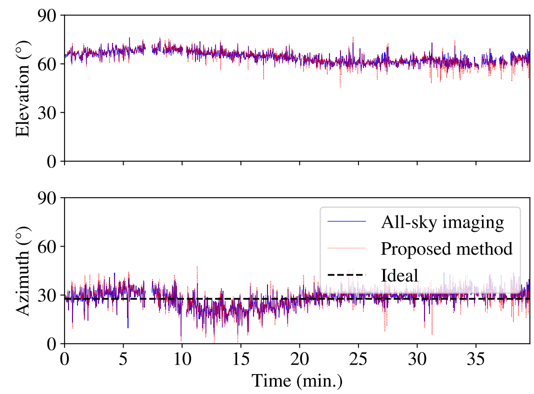

(b) 5.3515 MHz transmission frequency.

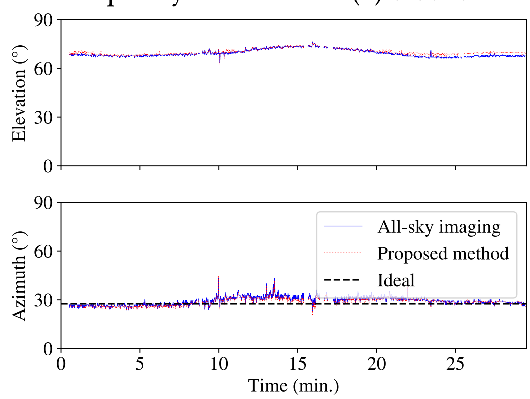

(c) 7.09 MHz transmission frequency.

Figure 3.7: DoA for signals transmitted by WK2XTU and received at LWA-SV using a 1 s integration time.

The visible decrease in DoA uncertainty at higher signal frequencies (fig. 3.7c versus fig. 3.7a) can likely be explained by the ionospheric channel. The highest frequency shown (7 MHz) is near the maximum usable frequency reflected by the ionosphere and thus there is likely only one signal path. Lower frequency signals generally travel multiple paths, which manifests as DoA uncertainty ("noise") in fig. 3.7. Figure 3.8 shows a histogram of azimuth values for our proposed method from the three observations shown in fig. 3.7 along with the geometric azimuth. This approximately Gaussian distribution with a mean at the geometric azimuth shows again that, despite channel effects, our DoA method functions as expected. The histograms are the same approximate shape independent of integration time but will have a varying sample count.

Processing a total of 6.9 h (24832 1 s integrations) of data yielded a consistently accurate result and showed that, when run using similar computational resources, our method

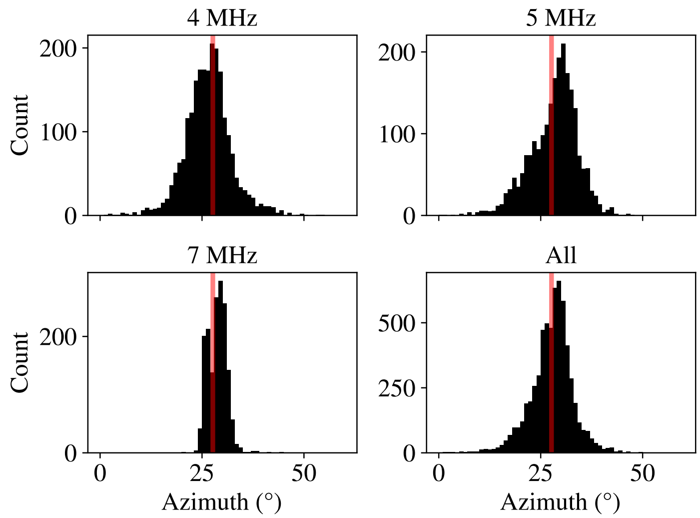

Figure 3.8: Histogram of azimuth values from our proposed method for the three observations shown in fig. 3.7. The geometric azimuth is shown in red. The last plot is the combination of the first three showing a frequency independent average.

was $\approx$ 18% faster than all-sky imaging.

3.4 Discussion

The results in fig. 3.7 show gradual changes in azimuth and elevation up to about $10°$ over time periods of tens of minutes. Under the assumption of specular reflection from a mirror ionosphere, these results may be interpreted as changes in the virtual height of the reflecting layer.

Results not shown using shorter integration times of 0.1 s and 0.5 s show no finer structure in the ionospheric virtual height measurements than is obtained with the 1 s integration time. A waterfall plot of the noise is approximately white independent of integration time.

The $\approx$ 18% speed-up from our method versus all-sky imaging is a rough estimate based on average computation-time-per-integration and over many hours of data $\approx$ 18% becomes substantial. For real-time processing—where the time to calculate DoA is less than the integration time—our method also offers more avenues for optimization than all-sky imaging does, since the fitting model and minimization algorithm can both be replaced.

We wish to highlight that our proposed method is more of a framework than a strictly defined algorithm. Our chosen minimization algorithm (Levenberg-Marquardt), which traverses the objective function, can be replaced with other traversal schemes. Similarly, our use of a wide 2-D Gaussian as a signal model when paired with the Levenberg-Marquardt removes the need for an initial estimate of signal DoA but can be replaced with any other signal model as befits the application. The accuracy, resolution, computational load, and processing time will vary depending on the chosen minimization algorithm and signal model.

We recognize that our use of visibility domain model fitting is not novel as radio astronomers regularly do more detailed variants of this. Likewise, we recognize that there exist sub-space based methods such as MUSIC for identifying signal DoA on arbitrary 2-D arrays. We present this method as an addition to the available toolkit for DoA estimation on truly arbitrary arrays.

The code for this work was written in Python and is available in a Git repository at `https://github.com/mistic-lab/lwa-tools`, which contains methods for processing the measured data received at LWA-SV as well as the simulation framework used in this chapter.

3.5 Conclusions

We have shown that our proposed method is more accurate than MUSIC and all-sky imaging for near-horizon narrowband signals, while being comparable for variations in SNR, array radius, frequency and frequency offset. The simulated result for low elevation signals is notable as our method provides a substantially more accurate DoA measurement near the horizon than MUSIC and all-sky imaging. Terrestrial signals refracted by the ionosphere travel across the Earth at low incidence angles, and our method can be used to accurately determine their source direction. We then show results from running both our method and all-sky imaging on several hours of a narrowband signal collected at LWA-SV. Processing of the real data showed no loss of accuracy down to the minimum integration time of 0.1 s. This shows the validity of our method outside of a simulated environment as the mean of the resulting azimuth matches the great-circle bearing from transmitter to the receiver.

Appendices

3.A Derivation of Complex Visibilities

Consider two antennas that are separated by a baseline vector $\vec{b}$. We can describe the baseline vector as a function of signal wavelength λ:

$$\vec{b} = b_x\hat{x} + b_y\hat{y} + b_z\hat{z} = \lambda\left(u\hat{x} + v\hat{y} + w\hat{z}\right) \tag{3.9}$$

Assuming a planar array, we neglect the $\hat{z}$ component. Let $E_1(l,m,t)$ and $E_2(l,m,t)$ denote the incident scalar electric field at the two antennas arriving from the direction indicated by the unit vector $\hat{s} = l\hat{x} + m\hat{y} + \sqrt{1-l^2-m^2}\hat{z}$ at time t. l and m, the $\hat{x}$ and $\hat{y}$ elements of $\hat{s}$, are called *direction cosines* (eq. (3.1)), and can be written as functions of elevation (θ measured from the horizon, $0°$, up to zenith, $90°$) and azimuth (ϕ measured from the local meridian, $0°$, clockwise while looking down). This is illustrated in fig. 3.2.

The intensity of the field incident on the array is

$$I(l,m,t) = \langle E_1(l,m,t)E_1^*(l,m,t)\rangle \tag{3.10}$$

where $\langle\cdot\rangle$ denotes expected value with respect to time. This value is estimated using the integral of the product of the signals over a finite integration time τ.

$$\hat{I}(l,m,t) = \frac{1}{\tau}\int_t^{t+\tau} E_1(l,m,x)E_1^*(l,m,x)dx \tag{3.11}$$

The separation of the two antennas means that the incident wave from $\hat{s}$ travels a distance of $\vec{b}\cdot\vec{s}$ between the two antennas, inducing a time delay in the signal received by the second. $E_2(l,m)$ is therefore

$$E_2(l,m,t) = E_1(l,m,t)e^{j\frac{2\pi}{\lambda}\vec{b}\cdot\vec{s}} \tag{3.12}$$

Since the two antennas are sensitive to radiation from the entire sky, the signals seen at their outputs are, respectively,

$$\begin{aligned} X_1(t) &= \iint E_1(l,m,t)\,dl\,dm \\ X_2(t) &= \iint E_2(l,m,t)\,dl\,dm \end{aligned} \tag{3.13}$$

The visibility associated with this pair of antennas is the correlation of their outputs expressed as a function of baseline coordinates

$$\begin{aligned} V(u, v, t) &= \langle X_1(t) X_2(t)^* \rangle \\ &= \iint \langle E_1(l, m, t) E_1^*(l, m, t) \rangle e^{j2\pi(ul+vm)}\, dl\, dm \\ &= \iint I(l, m, t) e^{j2\pi(ul+vm)}\, dl\, dm \end{aligned} \tag{3.14}$$

Exchanging the order of the expected value and integration requires that the source is spatially incoherent, which is trivially true for a point source. Every baseline b_k given by a pair of antennas results in a single sample $V(u_k, v_k)$ of the continuous visibility function. Equation (3.14) takes the form of a 2-D Fourier transform, and the incident radiation intensity can be recovered with an inverse Fourier transform.

$$\begin{aligned} I(l, m) &\xrightarrow{\mathcal{F}} V(u, v) \\ V(u, v) &\xrightarrow{\mathcal{F}^{-1}} I(l, m) \end{aligned} \tag{3.15}$$

To generate an all-sky image of $I(l, m)$, the inverse transform must be evaluated on a regular grid of u, v points. In general, this is computationally expensive, although resampling the visibility samples onto an appropriate grid permits the use of the FFT and a significant speedup. Visibility modelling, described above, avoids inverting this transform altogether.

It is also crucial to understand that the correlator output is integrated over time τ to estimate the expected value of the product in equation eq. (3.14). Whereas each antenna may be sampled at a high rate f_s, the visibilities are delivered at a lower rate $\frac{1}{\tau}$. In that time, the correlated values are repeatedly summed. In radio astronomy each output is called an *integration* whereas in MUSIC and similar algorithms it is known as a *snapshot*.

References

[1] R. Schmidt, "Multiple emitter location and signal parameter estimation," *IEEE Transactions on Antennas and Propagation*, vol. 34, pp. 276–280, Mar. 1986.

[2] H. Park, V. González-Gambau, and A. Camps, "High angular resolution RFI localization in synthetic aperture interferometric radiometers using direction-of-arrival

estimation," *IEEE Geoscience and Remote Sensing Letters*, vol. 12, pp. 102–106, Jan. 2015. Conference Name: IEEE Geoscience and Remote Sensing Letters.

[3] A. Barabell, "Improving the resolution performance of eigenstructure-based direction-finding algorithms," in *ICASSP '83. IEEE International Conference on Acoustics, Speech, and Signal Processing*, vol. 8, pp. 336–339, 1983.

[4] F.-G. Yan, J. Wang, S. Liu, Y. Shen, and M. Jin, "Reduced-complexity direction of arrival estimation using real-valued computation with arbitrary array configurations," *International Journal of Antennas and Propagation*, vol. 2018, pp. 1–10, 2018.

[5] R. Roy and T. Kailath, "Esprit-estimation of signal parameters via rotational invariance techniques," *IEEE Transactions on Acoustics, Speech, and Signal Processing*, vol. 37, no. 7, pp. 984–995, 1989.

[6] L. Tsung-Hsien and J. Mendel, "Azimuth and elevation direction finding using arbitrary array geometries," *IEEE Transactions on Signal Processing*, vol. 46, pp. 2061–2065, July 1998.

[7] O. A. Oumar, M. F. Siyau, and T. P. Sattar, "Comparison between MUSIC and ESPRIT direction of arrival estimation algorithms for wireless communication systems," in *The First International Conference on Future Generation Communication Technologies*, pp. 99–103, 2012.

[8] E. B. Daher, *Analysis and Design of Non-Uniform Arrays for Direction Finding*. PhD thesis, Villanova University, Villanova, Pensylvania, 2016. ISBN: 978-3-642-18058-3.

[9] E. Polisensky, "Station configuration studies writeup," Tech. Rep. 14, Long Wavelength Array, Dec. 2004.

[10] T. Torres, N. Anselmi, P. Nayeri, P. Rocca, and R. Haupt, "Low discrepancy sparse phased array antennas," *Sensors*, vol. 21, p. 7816, Nov. 2021.

[11] G. B. Taylor, C. L. Carilli, and R. A. Perley, *Synthesis Imaging in Radio Astronomy II*, vol. 180. Astronomical Society of the Pacific, 1999. Publication Title: Synthesis Imaging in Radio Astronomy II ISSN: 1050-3390.

[12] A. R. Thompson, J. M. Moran, and G. W. Swenson, *Interferometry and synthesis in radio astronomy*. Springer, 3 ed., 2017. arXiv: 1011.1669v3 ISSN: 09646639.

[13] J. F. Helmboldt, B. B. Markowski, D. J. Bonanno, T. E. Clarke, J. Dowell, B. C. Hicks, N. E. Kassim, and G. B. Taylor, "The Deployable Low-Band Ionosphere and Transient Experiment," *Radio Science*, vol. 56, July 2021.

[14] A. Malyali, M. Rivi, F. B. Abdalla, and J. D. McEwen, "Radio galaxy detection in the visibility domain," *Monthly Notices of the Royal Astronomical Society*, vol. 486, pp. 2695–2704, June 2019. arXiv: 1810.12930.

[15] J. A. Högbom, "Aperture synthesis with a non-regular distribution of interferometer baselines," *Astronomy and Astrophysics Supplement Series*, vol. 15, p. 417, June 1974.

[16] J. Dowell, D. Wood, K. Stovall, P. S. Ray, T. Clarke, and G. Taylor, "The Long Wavelength Array Software Library," *Journal of Astronomical Instrumentation*, vol. 01, p. 1250006, Dec. 2012. arXiv: 1209.1576.

Chapter 4

DDC Pool: Efficient Down-Conversion of Signals for RFI Monitoring

> I resolved to stop accumulating and begin the infinitely more serious and difficult task of wise distribution.
>
> Andrew Carnegie, Wealth

4.1 Introduction

RA observatories may suffer degradation in data quality due to the presence of RFI from terrestrial and satellite radio sources. Because RFI can mask or distort celestial signals, a reliable RFI monitoring system that can identify and characterize RFI is an essential component of an RA observatory.

Traditional RFI monitoring systems typically operate on integrated power spectra, losing the waveform-level information needed for advanced signal characterization. Access to the complex baseband time series of detected signals enables richer analysis and classification, supporting tasks such as modulation recognition and signal source identification.

When considering the RFI monitor architecture illustrated in fig. 4.1 in which the instantaneous bandwidth is very high (multiple GHz), it is not practical or economic to store a continuous stream of the wideband time series to disk. Such an approach produces massive datasets with low analytical value per storage unit, requiring costly post-processing to extract useful information. To address the challenges of data volume and signal characterization, we leverage the DDC as a means to efficiently extract narrowband signals from wideband data.

There are several DDC architectures [1]. However, to efficiently extract multiple signals in parallel we propose a collection of dynamically assignable synthesis filterbanks following the ideas in [2–5], which we call a "DDC Pool".

Figure 4.1 illustrates the proposed RFI monitor architecture, highlighting where the DDC Pool operates as a part of the signal processing chain. The voltage streams are squared to generate power spectra, which are fed into a live spectrogram display as well as a signal detector. In the context of Chapter 2, the power spectra are the integrated spectrometer data. Here, the unintegrated voltage spectra are also fed directly into the DDC Pool. The DDC Pool also receives signal metadata from a signal detector which operates on the integrated power spectra.

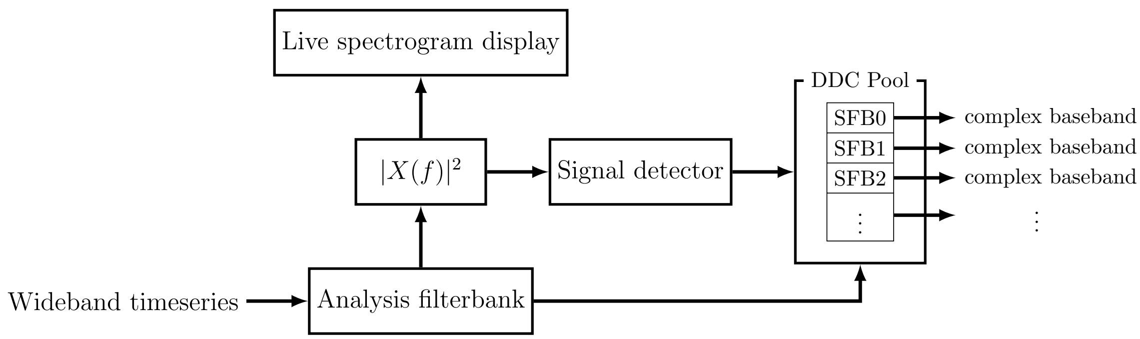

Figure 4.1: RFI monitor design concept block diagram.

The remainder of this chapter introduces the DDC Pool concept. Section 4.2 explains how the stages of the DDC Pool pipeline operate together along with demonstrations on a simple simulated dataset. Section 4.3 illustrates an application of the concept using a simulated wideband dataset which contains multiple signals of varying modulations. Section 4.4 concludes and discusses future work, while section 4.A outlines progress towards this goal.

4.2 Method

The proposed method reconstructs the detected signals' baseband waveforms from wideband data using three main stages:

1. channelization via an oversampled analysis filterbank,
2. signal detection in time-frequency space, and
3. synthesis of detected signals in a DDC Pool.

Rather than DDC architectures that process the entire wideband time series to down-convert each detected signal, the proposed method processes the wideband time series only once. The time series is channelized into oversampled channels using an analysis filterbank, then signals are identified in the time frequency spectrogram using a machine learning method. Using these detections and the oversampled channels, signals spread across multiple channels are synthesized to obtain their baseband waveforms.

To illustrate the method, we simulate a $BW = 100$ kHz wideband time series containing two frequency modulation (FM) signals of differing bandwidths. This simple dataset allows us to visualize how signals are split across filterbank channels and later reconstructed. Figure 4.2 shows the simulated data in the frequency domain along with analysis filterbank channels overlaid, which are discussed in section 4.2.1.

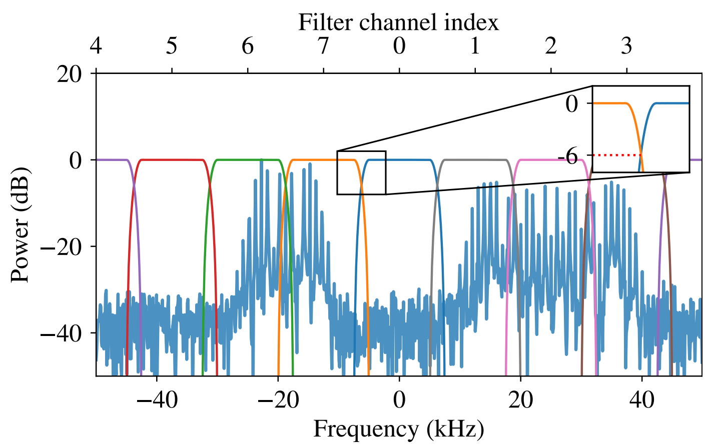

Figure 4.2: A simulated spectrum containing two FM signals and an overlaid oversampled filter showing the channels and -6 dB cross-over point.

4.2.1 Analysis Filterbank

The wideband time series is channelized in an oversampled polyphase filterbank (PFB) which we call the analysis filterbank. These channelized voltages are then squared to produce power spectra for signal detection and a live spectrogram display, as well as sent directly to the DDC Pool for detected signal synthesis.

The entire bandwidth containing multiple signals has a sample rate (f_s) equal to its bandwidth. This bandwidth is broken into M oversampled channels. Each channel has a bandwidth $BW = f_s/M$ proportional only to f_s and M. The channels are oversampled by a ratio (OS) greater than 1 and so have a channel sample rate proportional to the channel bandwidth as $f_{s,M} = OS \cdot BW$. An oversampling ratio of 2 is used as it is common in the literature. Figure 4.3 illustrates the relationship between these sample rates and bandwidths.

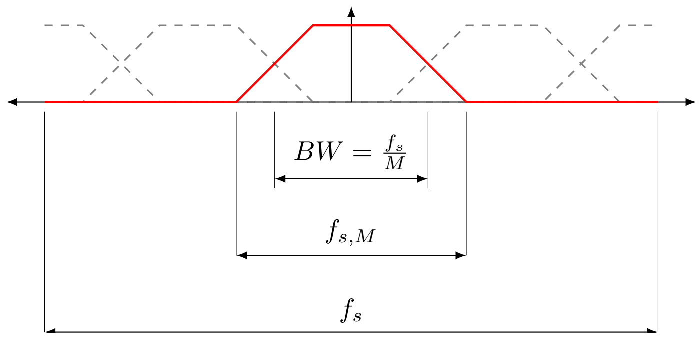

Figure 4.3: The relationship between the wideband time series sample rate (f_s), the bandwidth of each channel (BW), and the sample rate of each channel ($f_{s,M}$)

RF signals can span multiple frequency channels of the analysis filterbank, as shown in fig. 4.2.

A frequency-normalized polyphase raised-cosine filter with a roll-off factor (β) defined by the oversampling ratio (OS) is applied to each of the M analysis filterbank channels as described by eq. (4.1).

$$\begin{aligned} h(n) &= \text{sinc}(n)\frac{\cos(\pi\beta n)}{1-(2\beta n)^2} \\ n &= \left(\frac{-N}{2M}, \frac{N}{2M}\right) \end{aligned} \tag{4.1}$$

This ensures the oversampled channels cross at -6 dB (0.5 in linear units) so that signals can be reconstructed in voltage as described in [3].

The stopband attenuation is a function of the number of taps (N) in the filter as well as the transition bandwidth (Δf). The transition bandwidth is proportional to β via eq. (4.2). The channel bandwidth, BW, is equal to f_s/M, and is independent of the oversampling ratio.

$$\beta = \frac{2\Delta f}{BW} = \frac{2M\Delta f}{f_s} \tag{4.2}$$

The transition bandwidth is proportional to the oversampling ratio via eq. (4.3).

$$\Delta f = BW \cdot (OS - 1) \tag{4.3}$$

Realizing from eqs. (4.2) and (4.3) that the channel bandwidth is proportional to the sample rate by the oversampling ratio, $OS = f_s/BW$, we can express the transition width in terms of the sample rate and oversampling ratio exclusively as in eq. (4.4).

$$\Delta f = f_s \cdot \frac{OS - 1}{OS} \tag{4.4}$$

Two simulated FM signals are each split into multiple channels. The eight channels are shown in fig. 4.4 where the zeroth channel is centered over DC and the channels are indexed as they increase in frequency before wrapping to $-f_s/2$ and back up to DC. Note that there are eight channels, each with a sample rate of 25 kHz, spanning the 100 kHz bandwidth. This is because the filterbank oversampling ratio is 2.

4.2.2 Signal Detections

The output of the analysis filterbank is squared to produce power spectra. Signal detections are done on these power spectra, using the algorithm introduced in [6], in which an ML-based analysis algorithm outputs bounding boxes in time-frequency space. These bounding boxes are illustrated on the simulated data with red in fig. 4.5.

4.2.3 DDC Pool

The DDC Pool coordinates multiple synthesis filterbanks to reconstruct signals that span several channels, dynamically assigning computational resources based on detections.

As shown in fig. 4.1, the DDC Pool is fed with the bounding boxes detected in the power spectra (in which signal structure has been destroyed) along with the oversampled voltage channels from the analysis filterbank (in which signal structure is preserved). The DDC Pool handles the logic to pass only those voltage channels that overlap with a bounding box to a synthesis filterbank.

Synthesis Filterbank

Each synthesis filterbank reconstructs a single signal by aligning, padding, filtering, and combining the relevant analysis channels. The key operations are shown in fig. 4.6 and described below.

Each of the M_s analysis channels has a sample rate of $2f_s/M$, where M is the number of channels output from the analysis filterbank and f_s is the sample rate of the time series fed into the analysis filterbank input.

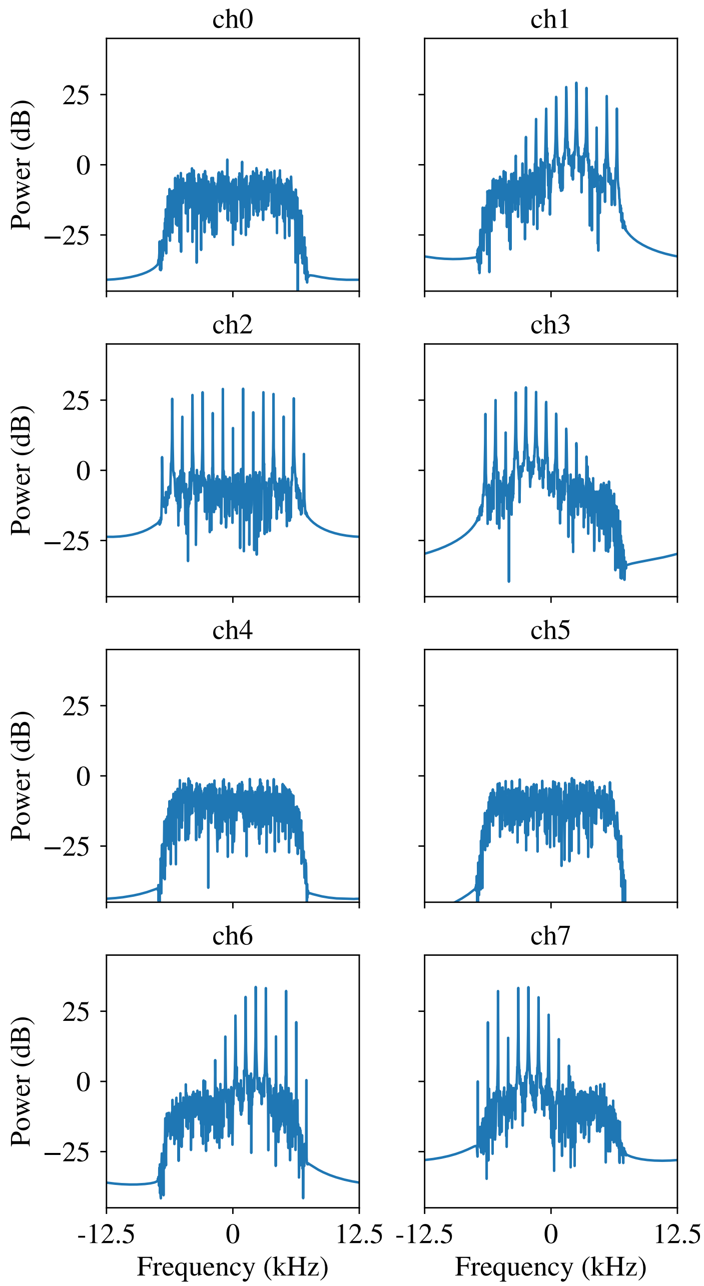

Figure 4.4: Spectra of $OS = 2$ oversampled channels output from the analysis filterbank.

The channels are fed into the synthesis filterbank in order of increasing frequency.

Padding

The padding process first checks whether M_s is even or odd. This defines the number of zero-valued channels used for padding: one is appended for an odd number of input channels while two are appended for an even number of input channels. These zero-

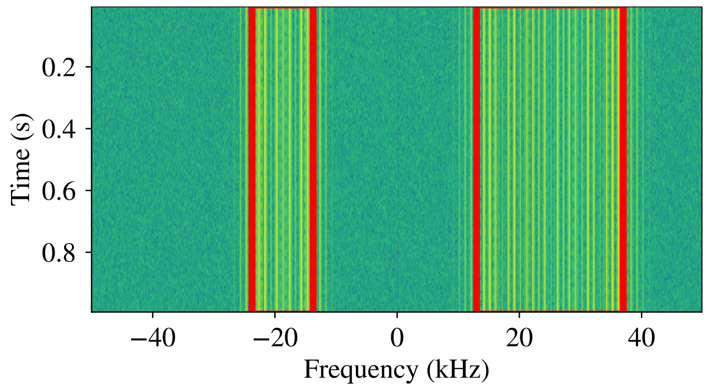

Figure 4.5: Detected bounding boxes (red) showing the frequency limits of the simulated FM signals.

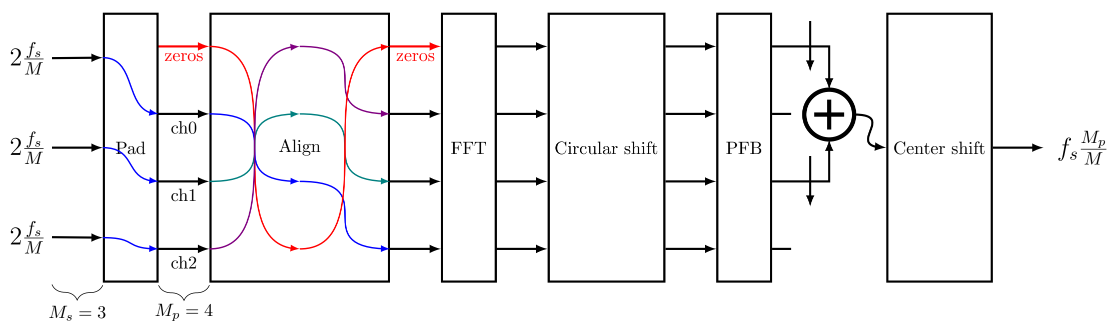

(a) A 3-channel case such that $M_s = 3$ and $M_p = 4$

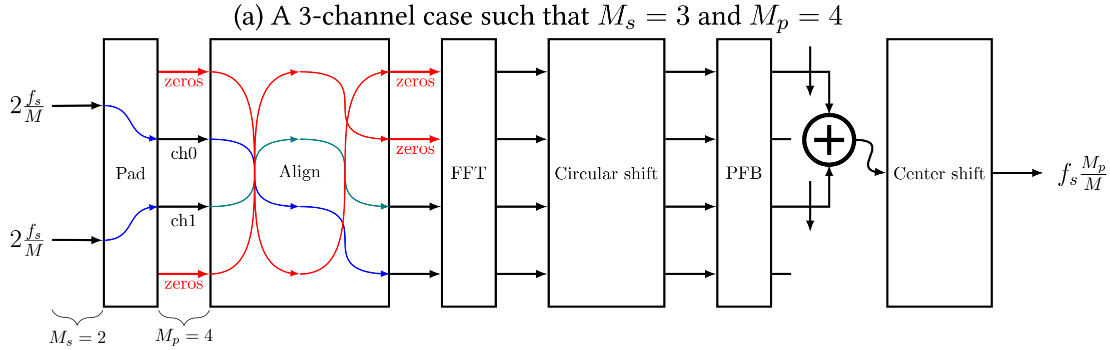

(b) A 2-channel case such that $M_s = 2$ and $M_p = 4$

Figure 4.6: Block diagram of synthesis filterbank processing steps. The M_s input channels are padded to an even number of channels, M_p, which are then synthesized into a single output time series.

valued padding channels ensure that the FFT and PFB always have an even channel count and that during the FFT the channel that wraps from $\frac{f_s}{2} \to -\frac{f_s}{2}$ is zeros, preventing any signal fragments from being wrapped to the wrong spectral locations.

At the output of the padding block there are M_p channels as decided by eq. (4.5).

$$M_p = \begin{cases} M_s + 2, & \text{where } M_s \bmod 2 = 0 \\ M_s + 1, & \text{where } M_s \bmod 2 \neq 0 \end{cases} \tag{4.5}$$

When M_s is odd, the single padding channel is inserted before the zeroth index. When M_s is even, padding channels are added before the zeroth index and after the M_sth index, effectively sandwiching the M_s input channels between two channels of zeros.

From this point onward, the processing steps within the synthesis filterbank will operate on M_p channels.

Align

The channels are aligned going into the FFT by first inverting their order then shifting them all by one index such that the channel containing the DC bin is kept in the center of the FFT.

If the dataset is static (i.e. not a stream of data) this alignment block can be replaced with a simple time reversal since the Fourier transform (FT) property of time reversal $x(-t) \leftrightarrow X(-f)$, or in this discrete case, the discrete Fourier transform (DFT) property of time reversal $x(N-n) \leftrightarrow X(N-k)$, shows that the two operations are equivalent. Of course, if the data are being streamed from a radio peripheral device or other data-generating input source, time reversal is not possible and the channel alignment must be done instead.

FFT

An M_p-point FFT is computed.

Circular Shift

The channels output from the FFT process are each phase rotated for alignment into the polyphase filter. Since the channels are oversampled, each must be shifted by $e^{-j2\pi kM/2}$ where k is the channel index. This shift brings the channel centers back to a spacing of f_s/M.

This can be accomplished by applying a complex rotator to each channel, or by applying the DFT time shift property, $x(n-m) = X(k)e^{-j2\pi km}$ such that each channel is simply delayed by $kM/2$ samples.

PFB

The next step is to apply the polyphase filter. The filter is a Nyquist filter (as in the analysis filterbank) with the same roll-off factor (β) but where $M = M_p/2$. Equation (4.1) can then be written as eq. (4.6).

$$
\begin{aligned}
h(n) &= \text{sinc}(n) \frac{\cos(\pi \beta n)}{1 - (2\beta n)^2} \\
n &= \left(\frac{-N}{M_p}, \frac{N}{M_p} \right)
\end{aligned}
\tag{4.6}
$$

As described in [7], the filter for each branch of the PFB is zero-stuffed and the second half of the channels have the filter taps delayed by 1 sample as illustrated by example in eq. (4.7) where $h(n)$ is the result of eq. (4.6) and $M_p = 4$.

$$
\begin{aligned}
h(n) =& [x_0, x_1, x_2, x_3, x_4, x_5, x_6, x_7] \\
h_0(n) =& [x_0, 0, x_4, 0] \\
h_1(n) =& [x_1, 0, x_5, 0] \\
h_2(n) =& [0, x_2, 0, x_6] \\
h_3(n) =& [0, x_3, 0, x_7]
\end{aligned}
\tag{4.7}
$$

Decimation

With the filter of eq. (4.6) applied, the output is summed over every $M_p/2$ samples, decimating the output sample rate to $f_s M_p / M$.

Center Shift

Finally, if M_s is odd the center frequency of the input bandwidth (the bandwidth covered by the M_s input channels) is frequency shifted to $0\,\text{Hz}$ on the synthesized spectrum. This is only necessary when M_s is odd; when M_s is even the “sandwiching” zero-valued padding channels do not cause an offset from the input center frequency.

When needed, the shift is proportional to M_p as $e^{j\pi n/M_p}$, where n is the sample index.

4.2.4 Post-Processing

Post-processing corrects residual frequency offsets and sample rate mismatches introduced during synthesis. Note that the final stage of the synthesis filterbank, the center shift, does not consider the center frequency of the signal, but of the input channels. The synthesis filterbank does not consider the concept of a detected signal and only synthesizes input channels. In the common case where a detected bounding box is not centered with respect to the analysis filterbank channels, a second shifting operation is required after the synthesis filterbank to shift the center of the detection box to $0\,\mathrm{Hz}$. Another effect of the zero-valued padding channels is that it changes the sample rate of the synthesized signal relative to the $OS = 2$ oversampled fragments. This can be handled by a simple decimation process if desired.

4.2.5 Output

Recall that fig. 4.2 shows one FM signal centered at $-18.75\,\mathrm{kHz}$ spanning 2 of the 8 channels and another centered at $25\,\mathrm{kHz}$ spanning 3 of the 8 channels.

Figures 4.7a and 4.7b show that both synthesized FM signals are reconstructed cleanly in baseband, validating the DDC Pool's operation. In the 3-channel case, the synthesis filterbank adds one zero-padding channel, while in the 2-channel case it adds two. The output sample rate considers these padding channels and so in both cases is $50\,\mathrm{kHz}$. This is because following eq. (4.5), in both cases ($M_s = 2$, $M_s = 3$) the total number of channels synthesized, including padding channels, is $M_p = 4$.

In order to increase the efficiency of the DDC Pool, two optimizations are made:

1. No synthesis filterbanks are pre-generated. If a detection box intersects some number of channels for which a synthesis filterbank has not yet been generated, a software object is created containing the appropriate zero-padding and channel alignments and then cached along with the filter taps specific to that channel count. This "deferred synthesis filterbank generation" results in a slower synthesis process the first time some number of channels is synthesized but a much faster process for signals of the same channel intersection count detected later.

   Note that the cached synthesis filterbank has no concept of the frequency for each channel it is synthesizing, and only cares about the number of channels being synthesized. For example, the first time a signal spans 3 channels, a synthesis filterbank is created. When another signal is detected that spans 3 channels the same synthesis filterbank can be used independent of any detection center frequencies.

2. If a known number of output samples is desired, the DDC Pool passes in twice that number of samples (the synthesis filterbank includes a $2\times$ decimation to return the signal to the input sample rate) plus the number of samples taken up by the filter transient. This ensures the minimum number of samples are processed instead of synthesizing the entire detected signal and then discarding output samples.

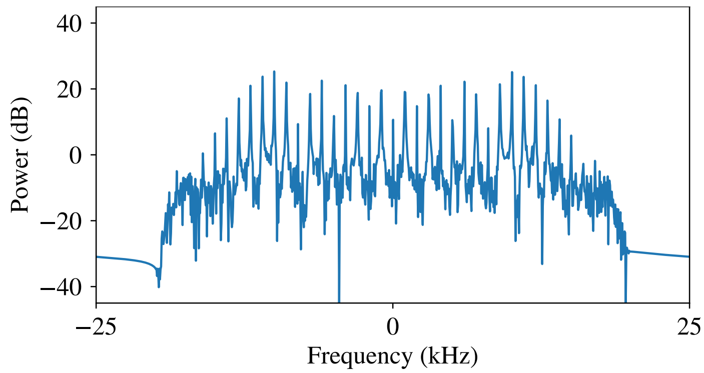

(a) Synthesized channels 1, 2 and 3.

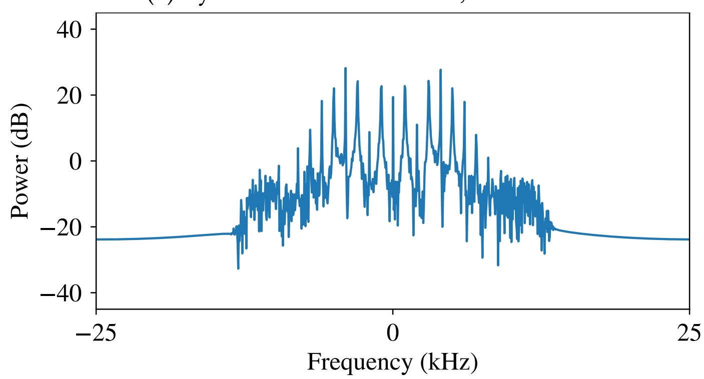

(b) Synthesized channels 6 and 7.

Figure 4.7: Synthesized FM signals.

4.3 Application

To demonstrate the practical utility of the DDC Pool, we apply it to a publicly available wideband dataset and synthesize short complex baseband segments suitable for ML classification tasks. We have found segments of 128-1024 complex time samples are suitable, although longer segments could be synthesized. In the following example, 256 synthesized samples are extracted for each detection.

Figure 4.8 shows the detected signals (top) and the synthesized baseband waveforms for several of them (bottom), confirming correct reconstruction across multiple modulation types, using a publicly available dataset [8].

The time series was passed through an analysis filterbank with $M_s = 16$, and the three signals indicated by non-red bounding boxes were synthesized using a series of synthesis filterbanks. The waveform synthesis can be validated visually by inspecting the symbols: the FSK2 (black) signal clearly switches between two frequencies, the PSK4 (orange) signal has four peaks at $1 + 0j$, $0 + 1j$, $-1 + 0j$ and $0 - 1j$, and the FSK4 (blue) signal switches

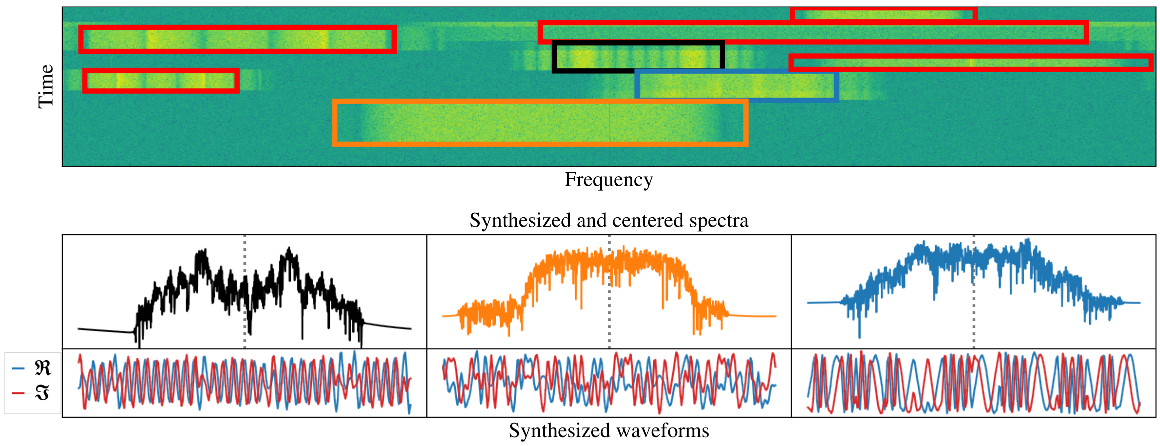

Figure 4.8: Waterfall plot (top) showing detection boxes with some of the signals synthesized (bottom)

between four frequencies. Table 4.1 shows details for each of the synthesis operations.

Table 4.1: Information on synthesized signals shown in fig. 4.8.

	Black	Orange	Blue
Number of synthesized channels	4	8	5
Number of padding channels	2	2	1
Modulation type	FSK2	PSK4	FSK4

Each synthesized signal has been centered with respect to the synthesis bandwidth following the steps in section 4.2.4. This corrects for the residual frequency offset between the analysis filterbank channels center frequencies and the detection box center frequency. There may still be a residual frequency offset caused by error in the bounding box detection algorithm.

4.4 Conclusions

This chapter introduced and validated the DDC Pool concept. It is a modular, computationally efficient architecture for extracting signal waveforms from wideband data using dynamically assignable synthesis filterbanks. The method has been validated on a publicly available wideband time series from which the complex baseband waveforms of detected signals were all extracted.

The application of this method to RFI monitoring is novel and suits the desire for a modular software-defined style of RFI monitoring. The method enables real-time RFI monitoring with reduced storage requirements while preserving the waveform information critical for advanced signal classification.

Future work will focus on generalizing the method to fractional oversampling ratios ($1 < OS < 2$), which would further reduce data volume while maintaining reconstruction fidelity. Progress made toward this goal is presented in section 4.A.

Appendices

4.A Fractional Oversampling Rates

4.A.1 Introduction

Chapter 4 introduced the DDC Pool concept, in which a wideband time series is channelized using an oversampled polyphase filterbank (analysis filterbank), signals are detected in the resulting time–frequency data, and the corresponding channels are recombined through a synthesis filterbank to reconstruct each signal's baseband waveform. The implementation in Chapter 4 assumed an oversampling ratio of two, ensuring 6 dB crossover points between adjacent channels and enabling straightforward reconstruction.

A natural extension of this work is to generalize the method to fractional oversampling ratios where $1 < OS < 2$. Fractional oversampling offers trade-off between data rate and reconstruction fidelity: it reduces the total data volume relative to $OS = 2$ while retaining sufficient overlap between channels for accurate synthesis. This efficiency is particularly valuable for wideband receivers, where doubling the sampling rate incurs significant computational and storage cost.

Although the literature discusses multirate and oversampled filterbanks [2–4, 9, 10], implementation details for fractionally oversampled synthesis filterbanks remain sparse. The goal of this appendix is therefore to formalize and clarify the design of fractionally oversampled analysis and synthesis filterbanks suitable for use within the DDC Pool architecture.

The sections that follow present this investigation in five parts. Section 4.A.2 derives a fractionally oversampled analysis filterbank based on the general multirate model. Section 4.A.3 introduces an ideal filter design that achieves perfect 0.5-linear crossover points for arbitrary oversampling ratios. Section 4.A.8 examines the synthesis filterbank and discusses current implementation challenges. Section 4.A.9 explores the interdependence of the analysis and synthesis filters, and section 4.A.10 outlines possible optimizations and directions for future work.

4.A.2 Analysis Filterbank

In this section, we extend the DDC Pool concept to fractional oversampling rates by deriving a general form for the analysis filterbank. This filterbank channelizes a wideband time series into partially overlapping channels suitable for subsequent synthesis. The derivation follows the general multirate filterbank model described in [10].

Let R be the channel sample rate reduction ratio, and M be the number of frequency channels. We are interested in the case where R and M are unrelated ($R \not\propto M$), with R still constrained to be an integer. However, if we have a desired OS ratio and a known/desired number of channels, M, we select R as the nearest integer that works, as $OS = M/R$.

Recall that $M/R = 1$ corresponds to a critically sampled filterbank, $R > M$ is an undersampled filterbank, $R < M$ is an oversampled filterbank. Here we are interested in oversampled filterbanks. In the $OS = 2$ oversampled case shown in Chapter 4, $R = M/2$ such that each channel is reduced from the wideband rate until it is twice what a critically sampled channel would be.

We begin with the general multirate model for an analysis filterbank, expressed as

$$X_m(k) = \sum_{n=-\infty}^{\infty} h(kR - n)x(n)W_M^{-mn}, \quad m = 0, 1, \ldots, M-1\,, \tag{4.8}$$

where $h(n)$ is the analysis filter[1], $x(n)$ is the input time series, m is the frequency channel, n is the input sample number in time, k is the output sample number in time, sometimes called the decimated sample time, which also denotes the block number[2], and $W_M^{-mn} = e^{-j2\pi nm/M}$, the complex modulation function. The output, $X_m(k)$, is the short-time spectrum of the mth channel at decimated time index $k = n/R$,

$$X_m(k) = \sum_{n=-\infty}^{\infty} y_k(n)W_M^{-mn}\,, \tag{4.9}$$

where

$$y_k(n) = h(kR - n)x(n)\,. \tag{4.10}$$

In this formulation, $h(n)$ is a sliding analysis window. The length of $h(n)$ and the number of samples in the DFT, M, determine the output time/frequency resolution.

Sliding Versus Stationary Time Frame

Equation (4.9) assumes stationary data with a sliding analysis window. In a real-time system in which data arrives continuously, we must have a fixed filter while the input stream advances. We accomplish this by transforming to a sliding time frame using the

[1] An ideal analysis filter is introduced below in section 4.A.3.

[2] A block consists of one sample for each output frequency channel

substitution $r = n - kR$, where r tracks time relative to the current block. The filterbank equation then becomes

$$X_m(k) = \sum_{r=-\infty}^{\infty} h(-r)x(r+kR)W_M^{-m(r+kR)} \tag{4.11}$$

$$= W_M^{-mkR} \underset{\sim}{X}_m(k) \tag{4.12}$$

where

$$\underset{\sim}{X}_m(k) = \sum_{r=-\infty}^{\infty} h(-r)x(r+kR)W_M^{-mr} \tag{4.13}$$

is the short-time Fourier transform (STFT) with respect to the origin of the sliding time frame[3]. This form aligns with the desired real-time implementation.

Equation (4.13) is not in the correct form for an M-point FFT because summing over r may span more than M samples. To convert to a form appropriate for an M-point FFT, we time-alias the signal

$$\underset{\sim}{y}_k(r) = h(-r)x(r+kR) \tag{4.14}$$

such that eq. (4.13) becomes

$$\underset{\sim}{X}_m(k) = \sum_{r=-\infty}^{\infty} \underset{\sim}{y}_k(r) W_M^{-mr} \tag{4.15}$$

where $\underset{\sim}{y}_k(r)$ can be time aliased into an M-point sequence. The time-aliased M-point sequence, $\underset{\sim}{x}_k(r)$, is constructed by subdividing $\underset{\sim}{y}_k(r)$ into blocks of M samples and then stacking and adding the blocks:

$$\begin{aligned} \underset{\sim}{x}_k(r) &= \sum_{l=-\infty}^{\infty} \underset{\sim}{y}_k(r+lm) \\ &= \sum_{l=-\infty}^{\infty} h(-r-lm)x(r+lm+kR), \quad r = 0, 1, \ldots, M-1\,. \end{aligned} \tag{4.16}$$

Now the STFT can be written as

$$\underset{\sim}{X}_m(k) = \sum_{r=0}^{M-1} \underset{\sim}{x}_k(r) W_M^{-mr} \tag{4.17}$$

[3] $\sim$ under a variable denotes a signal in the sliding time frame

which works for an M-point FFT.

In summary, the fractional analysis filterbank can be implemented by loading R new samples per block into a shift register of length equal to the filter's number of taps, weighting them by $h(-r)$, time-aliasing into M-sample blocks, and computing an M-point DFT to obtain $\underset{\sim}{X}_m(k)$. This process generalizes the integer oversampling case and supports arbitrary fractional oversampling ratios.

We have implemented this method in Python and verified correct operation for multiple fractional values of OS, providing a flexible foundation for the fractionally oversampled DDC Pool.

4.A.3 Ideal Analysis Filterbank Coefficients

The previous section derived a general fractional oversampling PFB. In this section we define an ideal filter response for such filterbanks. An ideal filter guarantees perfect 0.5-linear unit crossover points, and overlapping transition widths that sum to unity, between adjacent channels for any fractional oversampling ratio. This design ensures a flat passband when overlapping channels are summed, enabling perfect waveform reconstruction in the synthesis stage.

Filter Definition

The filter is a frequency-normalized polyphase raised-cosine filter. The impulse response is

$$h(n) = \text{sinc}(n)\frac{\cos(\pi\beta n)}{1-(2\beta n)^2}, \quad n \in \left(\frac{-N_h}{2K}, \frac{N_h}{2K}\right) \tag{4.18}$$

where β is the roll-off factor and N_h is the number of filter taps. The raised-cosine term provides smooth transition regions between channels. Frequency normalization ensures that the filter shape remains consistent across different sampling rates and numbers of channels, allowing the filter to scale with receiver bandwidth.

A key property of this filter is that the oversampled channels cross at -6 dB (0.5 in linear units) so that signals can be reconstructed in voltage as described in [3].

4.A.4 Parameter Relationships

The roll-off factor β controls the transition bandwidth Δf relative to the channel bandwidth BW:

$$\beta = \frac{2\Delta f}{BW} \tag{4.19}$$

and because the transition bandwidth scales with the oversampling ratio,

$$\Delta f = \frac{BW}{2} \times (OS - 1)\,. \tag{4.20}$$

Combining these expressions yields

$$\beta = OS - 1 = \frac{M}{R} - 1 \tag{4.21}$$

so that β can be determined directly from the fractional oversampling parameters M and R.

4.A.5 Filter Behaviour in Fractional Oversampling

Recall from section 4.2.1 that the total system bandwidth is equal to the wideband sample rate, f_s. The bandwidth of each channel is therefore $BW = f_s/M$, and the channels are oversampled by a ratio $OS > 1$. The sample rate for each channel is then

$$f_{s,M} = OS \times BW = \frac{M}{R} \times \frac{f_s}{M} = \frac{f_s}{R} \tag{4.22}$$

The transition bandwidth Δf is proportional to $(OS - 1)$, determining the degree of channel overlap. Fractional oversampling ratios between one and two provide a tunable balance between data rate and the robustness gained from higher channel overlap.

4.A.6 Final Filter Expression

Substituting the relationships above into eq. (4.18) yields a final expression for the ideal filter coefficients:

$$h(n) = \text{sinc}(n)\,\frac{\cos\left(2\pi \frac{M}{R} n^2\right)}{1 - \left(4n\left(\frac{M}{R} - 1\right)\right)^2}, \quad n \in \left(\frac{-N_h}{2M}, \frac{N_h}{2M}\right)\,. \tag{4.23}$$

This formulation depends only on the oversampling parameters M and R, making it applicable to any fractional oversampling ratio. Adjacent channels intersect at 0.5 linear amplitude (-6 dB), guaranteeing a flat combined response and perfect reconstruction.

4.A.7 Filter Validation

Figure 4.9 illustrates the frequency response of this filter for several fractional oversampling ratios. Each curve exhibits symmetrical transition bands and identical crossover points, confirming that the design maintains perfect-reconstruction behaviour across all tested values.

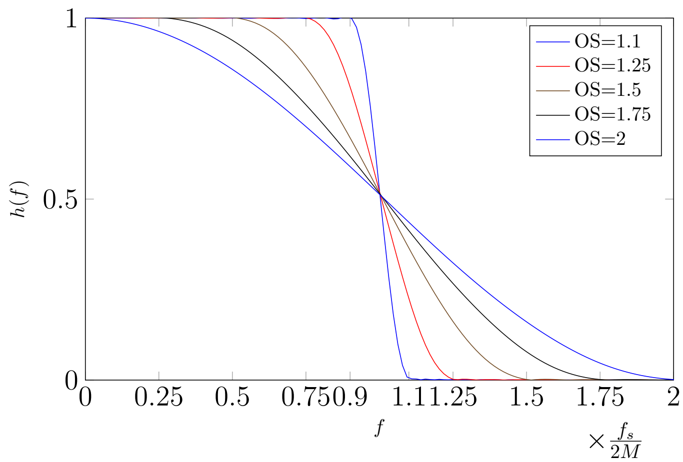

Figure 4.9: Frequency response of the ideal analysis filterbank filter introduced in eq. (4.23) for a variety of fractional oversampling ratios. Each channel exhibits a 0.5-linear (-6 dB) crossover point at the edge of its bandwidth, ensuring perfect reconstruction when neighbouring channels are combined.

We have implemented this filter in Python and verified its performance for both integer and fractional oversampling ratios. It serves as the ideal analysis filter for the fractional synthesis filterbank introduced in section 4.A.8.

4.A.8 Synthesis Filterbank

The synthesis filterbank is the dual of the analysis filterbank described in section 4.A.2. Its purpose is to reconstruct the original time-domain waveform, $\hat{x}(n)$, from the channelized spectra, $\hat{X}_m(k)$, produced by the analysis stage. In the context of the DDC Pool, the synthesis filterbank combines the oversampled channels that contain fragments of a detected signal to recreate its complex baseband waveform.

Derivation

Following [10], the basic structure of a synthesis filterbank can be written as

$$\hat{x}(n) = \sum_{k=-\infty}^{\infty} f(n-kR)\,\frac{1}{M}\sum_{m=0}^{M-1} \hat{X}_m(k)\,W_M^{mn} \tag{4.24}$$

where $f(n)$ is the synthesis filter impulse response, R is the decimation factor, M is the number of channels, and $W_M^{mn} = e^{j2\pi mn/M}$. The inner summation performs an inverse DFT across all subband signals, while the outer summation applies time-shifted versions of the synthesis filter to reconstruct the continuous time series.

Transformation to the Sliding Time Frame

To align with the analysis filterbank formulation, we express the synthesis operation in the sliding-time frame by substituting $r = n - k_0 R$, where k_0 denotes the current block index. Substituting into eq. (4.24) gives

$$\hat{x}(r+kR)\big|_{k=k_0} = f(r)\,\frac{1}{M}\sum_{m=0}^{M-1} \hat{X}_m(k)\,W_M^{mkR}\,W_M^{mr}\big|_{k=k_0} + \text{terms for } k \neq k_0\,. \tag{4.25}$$

The first term in eq. (4.25) represents the contribution from the current block k_0, while the remaining terms correspond to overlapping contributions from other blocks due to the finite filter length. These cross-block interactions are what ensure continuity and perfect reconstruction in the overall system.

Rewriting in Terms of a Sliding-Frame

Analogous to the analysis filterbank, we define the sliding-frame spectrum as

$$\underset{\sim}{\hat{X}}_m(k) = \hat{X}_m(k)\,W_M^{mkR} \tag{4.26}$$

and its inverse DFT as

$$\underset{\sim}{\hat{x}}_k(r) = \frac{1}{M}\sum_{m=0}^{M-1} \underset{\sim}{\hat{X}}_m(k)\,W_M^{mr}\,. \tag{4.27}$$

Substituting these definitions into eq. (4.25) yields

$$\hat{x}(r+kR)\big|_{k=k_0} = f(r)\,\underset{\sim}{\hat{x}}_k(r)\big|_{k=k_0} + \text{terms for } k \neq k_0\,. \tag{4.28}$$

This equation shows that the reconstructed signal at each block is the product of the synthesis filter and the inverse-transformed channelized data, plus additional contributions from adjacent blocks.

Implementation Challenges

If the filter length is N_f, only N_f/R samples contribute to each output sample $\hat{x}(n)$. When $N_f > M$, $\hat{\underset{\sim}{x}}_k(r)$ must be interpreted as a periodically extended sequence with period M. The literature indicates that the remaining terms in eq. (4.25) (for $k \neq k_0$) account for inter-block overlap and continuity [10]. However, the explicit implementation for fractionally oversampled systems is not clearly described.

In practice, this ambiguity means that while the theoretical form of the synthesis filterbank is known, a complete block-wise implementation that supports arbitrary fractional oversampling ratios has not yet been realized. The difficulty lies in managing the contributions from overlapping filter segments.

Discussion

The derivation above establishes the mathematical structure of the fractionally oversampled synthesis filterbank and its relationship to the analysis filterbank. In principle, perfect reconstruction requires that the analysis and synthesis filters satisfy the condition

$$\sum_{k=-\infty}^{\infty} f(n-kR)\, h(kR-n+sM) = \begin{cases} 1, & s=0 \\ 0, & \text{otherwise} \end{cases}, \tag{4.29}$$

ensuring that all aliasing terms cancel. For integer oversampling ratios (e.g., $OS = 2$), this condition can be satisfied using the raised-cosine filters introduced earlier. For fractional ratios, however, the interaction between the non-integer decimation factor and the overlapping filter segments introduces additional complexity.

Further work is therefore required to define a synthesis filterbank that accounts for the inter-block terms and achieves perfect reconstruction for arbitrary $1 < OS < 2$. Despite this limitation, the analysis presented here clarifies the mathematical requirements for a fractional synthesis filterbank and identifies where the current literature is incomplete.

4.A.9 Filter Interdependence

Perfect reconstruction in a filterbank system requires that the analysis and synthesis filters form a complementary pair. The filter relationship between the two was defined in eq. (4.29), where $h(n)$ and $f(n)$ are the analysis and synthesis filters respectively. This condition ensures that all aliasing terms cancel and that the reconstructed signal $\hat{x}(n)$ is identical to the original input $x(n)$.

For the integer oversampling case ($OS = 2$), $f(n)$ can be obtained directly as the time-reversed counterpart of $h(n)$ using the raised-cosine form defined in section 4.A.3. However, for fractional oversampling ratios ($1 < OS < 2$), the relationship between $h(n)$ and $f(n)$ becomes non-trivial because the decimation factor R and the number of channels M are not integer multiples. Deriving $f(n)$ that satisfies this condition for arbitrary fractional values of OS remains an open area of investigation.

4.A.10 Optimizations

Several optimizations reduce the computational complexity of the fractional DDC Pool architecture without compromising reconstruction accuracy. The following subsections outline potential improvements for the analysis filterbank, the synthesis filterbank, and the combined analysis/synthesis pair.

Analysis Filterbank

In the fractional analysis filterbank, the modulation term W_M^{-mkR}, which converts the stationary time-series representation to the sliding-time frame, can be implemented more efficiently. Instead of explicitly performing complex multiplications by W_M^{-mkR}, an equivalent operation can be achieved through a modulo-M circular shift of the input sequence $\underset{\sim}{x}_k(r)$ prior to the DFT. This equivalence is illustrated in fig. 4.10 and removes the need for a complex exponential computation per channel, thereby reducing per-frame processing cost.

Synthesis Filterbank

A similar optimization applies to the synthesis filterbank. The modulation term W_M^{kR}, which re-aligns the channelized data in the time domain, can also be implemented as a modulo-M rotation of the DFT output rather than an explicit complex multiplication.

Analysis/Synthesis Filterbank Pair

When the synthesis filterbank directly follows the analysis filterbank, as is the case within the proposed DDC Pool, the phase modulation stages in each can effectively cancel. Be-

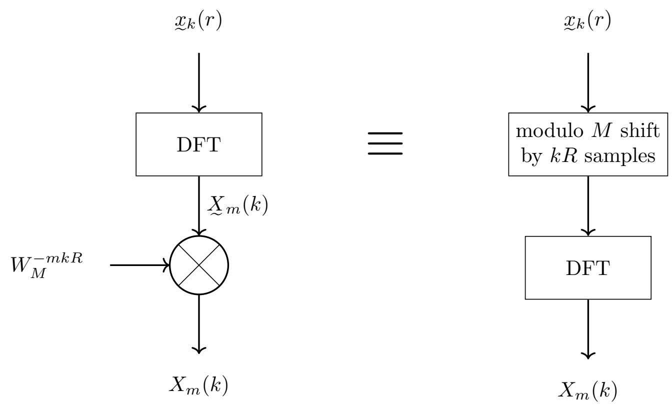

Figure 4.10: Modulo equivalence in the analysis filterbank.

cause the phase shift introduced by the analysis bank is equal and opposite to that of the synthesis bank, the intermediate modulation steps can be omitted without affecting the final reconstructed signal. This simplification is valid as long as no non-commutable operations are applied between the two stages (for example, non-linear processing or arbitrary resampling). Eliminating redundant modulation stages significantly reduces computational load and simplifies the implementation.

Summary

The optimizations described above reduce the number of complex multiplications required in both the analysis and synthesis filterbanks while preserving the perfect reconstruction properties of the fractional DDC Pool. Future work will implement and benchmark these optimizations within the DDC Pool architecture.

References

[1] A. Abinaya and M. Maheswari, "A survey of digital down converter architecture for next generation wireless applications," *IOP Conference Series: Materials Science and Engineering*, vol. 872, p. 012037, June 2020.

[2] F. Harris, C. Dick, X. Chen, and E. Venosa, "Wideband 160-channel polyphase filter bank cable TV channeliser," *IET Signal Processing*, vol. 5, no. 3, p. 325, 2011.

[3] F. Harris, E. Venosa, X. Chen, and B. D. Rao, "Polyphase analysis filter bank downconverts unequal channel bandwidths with arbitrary center frequencies," *Analog Integrated Circuits and Signal Processing*, vol. 71, no. 3, pp. 481–494, 2012.

[4] E. Venosa, F. J. Harris, and F. A. Palmieri, *Software Radio: Sampling Rate Selection, Design and Synchronization*, vol. 53. Springer, 2013. Publication Title: Analog Circuits and Signal Processing Issue: 9.

[5] T. W. Rondeau, T. J. O'Shea, and V. Tech, "Designing analysis and synthesis filterbanks in GNU Radio," in *Karlsruhe Institute of Technology's Workshop on Software Radio (WSR)*, (Karlsruhe), p. 9, 2014.

[6] S. Harrison, R. Coles, T. Robishaw, and D. Del Rizzo, "RFI novelty detection using machine learning techniques," in *RFI 2019*, 2019.

[7] D. Zhou, "A review of polyphase filter banks and their application," In-House Final Technical Report AFRL-IF-RS-TR-2006-277, Air Force Research Laboratory, Rome, New York, 2006.

[8] N. West, "SPAWC21 Wideband Dataset," May 2021. Type: dataset.

[9] F. Harris, *Multirate Signal Processing for Communication Systems.* S.l.: RIVER PUBLISHERS, 2021. OCLC: 1229088160.

[10] R. E. Crochiere and L. R. Rabiner, *Multirate digital signal processing.* Prentice-Hall signal processing series, Englewood Cliffs, N.J: Prentice-Hall, 1983.

Chapter 5

SigCLR: A Contrastive Learning Approach to Unsupervised Modulation Recognition and Novelty Detection

> The first principle is that you must not fool yourself—and you are the easiest person to fool.
>
> Richard P. Feynman, 1974

5.1 Introduction

As radio environments become more crowded, and RFI monitoring expands to wider bandwidths, a significant amount of RFI data is being collected. At the DRAO, efforts have evolved from building a wideband RFI monitor (Chapter 2), to detecting signals in wideband spectra [1], and extracting baseband time series of these signals (Chapter 4) while discarding the rest of the wideband spectra. This process reduces the data volume and recovers time-domain characteristics of the interfering signals. Such characteristics are lost in other common stored data types such as integrated power spectra and occupancy.

This data collection method supports rich RFI science, such as RF site characterization and signal identification, without storing large data volumes. However, the lack of labels in this dataset complicates tasks such as modulation recognition, feature vector extraction, and novelty detection.

We demonstrate a method for modulation recognition on an unlabelled dataset containing 53 modulations, generated using TorchSig [2]. Contrastive learning [3] is a self-supervised approach that builds useful representations without labelled data but benefits from large training sets; TorchSig's ability to generate arbitrarily large synthetic datasets makes it well suited to this approach.

We introduce SigCLR, a model architecture that adapts image-based contrastive learning to complex-valued radio time-series data from the TorchSig dataset.

5.2 Overview

Figure 5.1 illustrates the model's architecture and training process. Each time-series signal $s(t)$ is duplicated and augmented differently using RF-specific augmentations, a_i and a_j (e.g. gain drift, clipping, spectral inversion) from the same family $\mathcal{A}$. The augmented copies $s_1(t)$ and $s_2(t)$ are processed through identical encoders and projection heads. Their representations z_1 and z_2 are compared in a cross-entropy loss function to guide the network towards treating augmented views of the same signal as equivalent representations. After training, the augmentation pipeline and projection head are removed, leaving the encoder for downstream tasks. We demonstrate the encoder's ability to cluster signal types, effectively serving as a modulation recognizer.

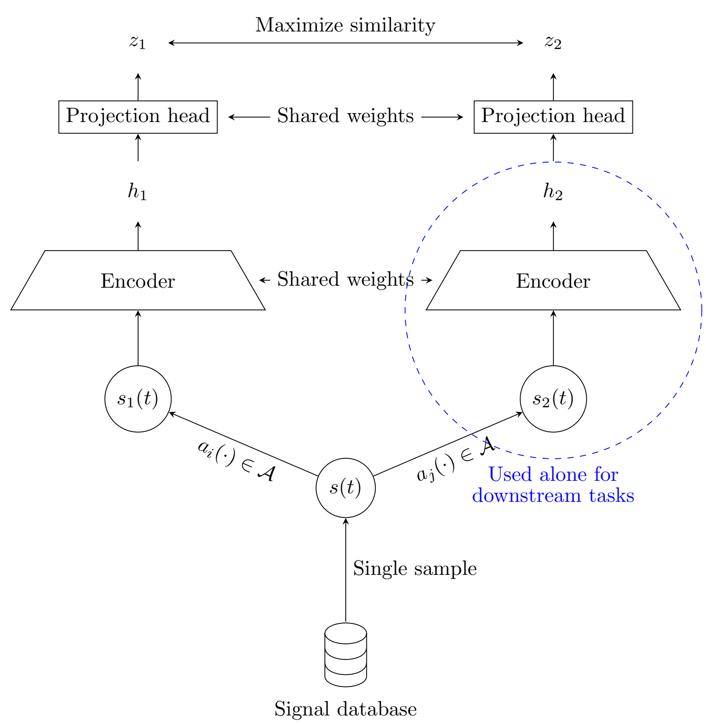

Figure 5.1: SigCLR model architecture.

5.3 Architecture

5.3.1 Augmentations

Each copy of the input time-series signal undergoes a single, randomly chosen augmentation from the following:

- Time-varying additive white Gaussian noise, in which up to 10 random regions have noise in the range of -80 to -20 dB applied.
- Random phase shift, in which the signal is shifted by a phase in the range of $-\pi$ to π.
- Time reversal, in which the sample order is inverted. In the frequency domain, this corresponds to conjugating the spectrum.
- Random time shift, in which the input is zero-padded to retain the input size.
- Gain drift, in which the input signal is modified as though the front end gain controller drifted as a random walk during the signal collection.
- Local oscillator drift, in which the input signal is modified as though the local oscillator drifted as a random walk in frequency during the signal collection.
- Clipping, in which the signal is clipped by a random percent in the range of 75% to 95%.
- Spectral inversion, in which the frequency sense of the signal is reversed. This corresponds to conjugating the signal in the time domain.

Figure 5.2 shows an exaggeration of these augmentations being applied to a 2GMSK modulated snippet from the training set.

5.3.2 Encoder

We use an EfficientNet B4 encoder [4] adapted from image to complex time series data [2]. The encoder classification layer is removed and a neck is added which customizes the encoder for a specific downstream task. In the case of unsupervised classification, the neck consists of:

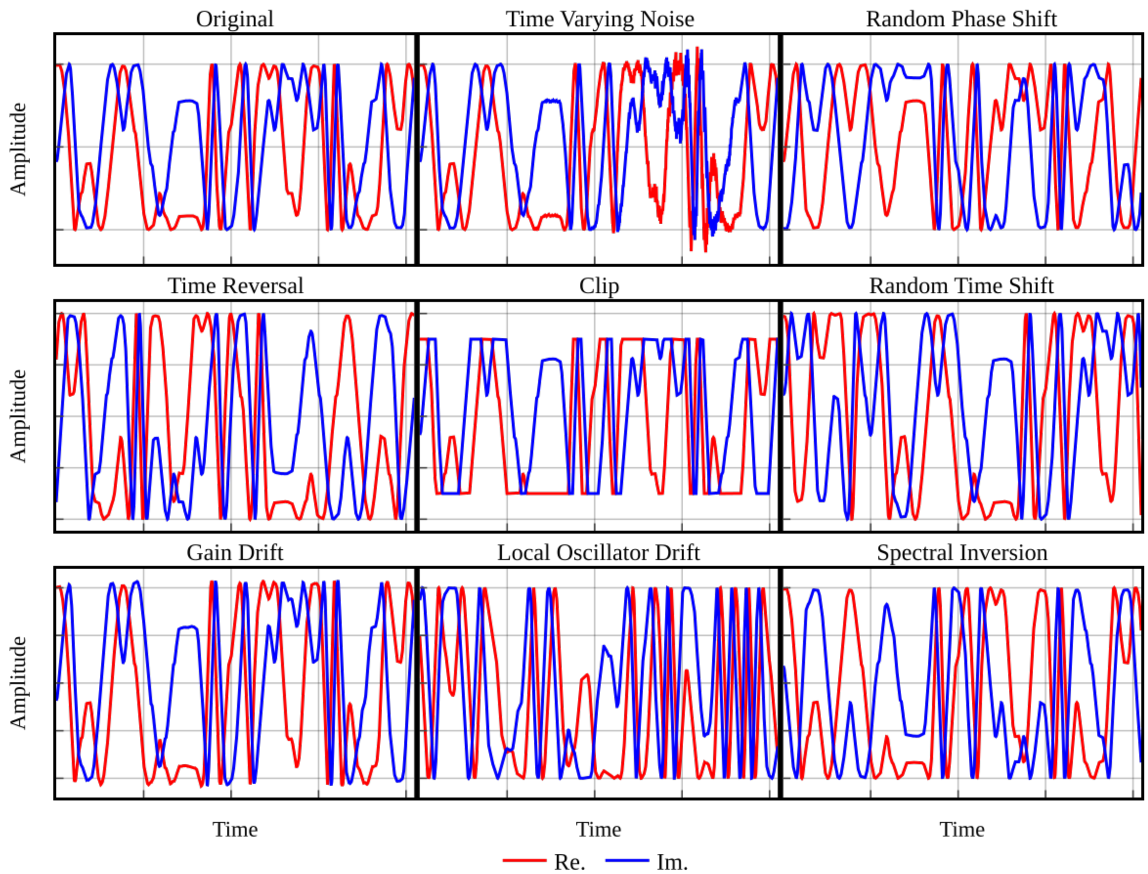

Figure 5.2: Augmentations applied to a 2GMSK signal. Some augmentations are exaggerated for visibility.

- a linear layer that expands the encoder output to 512 features, providing the opportunity to adapt the encoder to various downstream tasks,
- a Sigmoid Linear Unit function, chosen because of its known effectiveness in deep learning architectures,
- a final linear layer with an output size matching the number of desired classes.

In the case of unsupervised classification, we use the narrowband TorchSig dataset, which consists of 53 classes and so the final linear layer in the neck has an output dimension of 53.

5.3.3 Projection Head

A projection head is used during the training process to decouple the latent representation space from the task of contrastive learning. The projection head is a small neural network made up of:

- a linear layer with input and output sizes equal to the number of desired classes,
- a batch normalization,
- a Sigmoid Linear Unit function,
- an unbiased linear layer with input and output sizes equal to the number of desired classes.

Adding the projection head prevents the encoder from optimizing for the contrastive learning task—the encoder can then focus on learning high-quality signal representations while the projection head handles the contrastive objective.

After training, the projection head is discarded and the encoder is used alone for downstream tasks.

5.3.4 Similarity and Loss Metrics

We use cosine similarity to measure similarity between two signal representations, z_i and z_j, at the output of the projection head:

$$sim(z_i, z_j) = \frac{1}{\tau} cos(\theta) = \frac{1}{\tau} \frac{z_i \cdot z_j}{||x_i||_2 \cdot ||x_j||_2}. \tag{5.1}$$

Here, the L2 norm, $||x||_2$, computes the magnitudes, and τ (set to 0.07) acts as a temperature to penalize differing vectors and favour similar ones. Following [3], we use cross-entropy loss on the similarities with summing reduction as the loss criterion.

5.3.5 Optimizer and Hyperparameter Tuning

We use the AdamW optimizer with an initial learning rate of 0.001. To help prevent optimizing for local minima, we use a cosine annealing learning rate scheduler.

5.4 Batch Size and Training Process

The training is spread across three compute nodes, each with four $32\,\text{GB}$ NVIDIA V100 Volta GPUs. Each node sends a batch size of 64 to each GPU for a total effective batch size of 768. Training lasts for ≈ 800 epochs.

5.5 Method

Using TorchSig, we generate a training dataset of 530,000 unique signals evenly split across 53 modulation classes. Each signal has 4096 complex time series samples. The validation dataset contains 106,000 unique signals with a uniform class distribution.

We experimented with a LARS optimizer but did not see a performance gain. This was likely because of our relatively small batch size.

5.6 Results

To validate the training, we use the trained encoder as a modulation recognizer. Although the training was unsupervised, TorchSig provides labelled data. We generate two signals for each of the 53 modulations, then compute the cosine similarity between each pair without temperature normalization, visualized in a non-symmetric similarity matrix (see fig. 5.3). The asymmetry arises because signals are divided into two views, "A" and "B", and similarities are computed between them. These results demonstrate unsupervised modulation recognition, though performance is limited by the small dataset size.

We then group modulations: frequency-modulated signals under FSK, phase-modulated under PSK, amplitude-modulated under ASK, and OFDM signals kept distinct. The average similarity of each group is shown in fig. 5.4a. The simplified modulation recognizer performs better (fig. 5.4b), especially when OFDM is excluded, since OFDM is often confused with amplitude-modulated signals.

5.7 Conclusion

We demonstrate a method for training large models on unsupervised radio data using a contrastive learning architecture to build a simple modulation recognizer. We introduce the architecture and show results for classifying unknown signals as frequency, phase, or amplitude modulated.

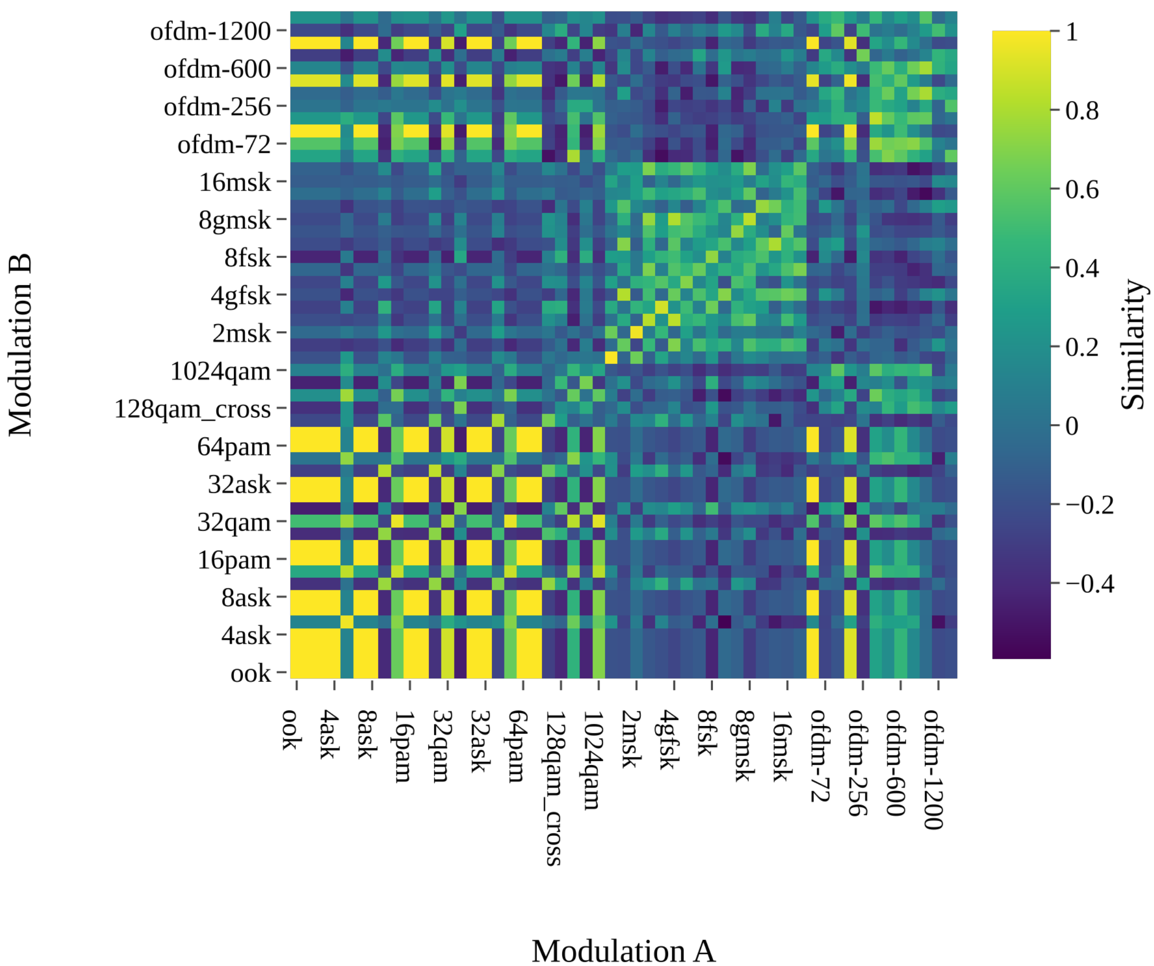

Figure 5.3: Similarity matrix for every modulation showing a trend towards the ideal diagonal. The most confused blocks are often similar modulations such as 8ASK and 4ASK. In general amplitude and phase modulations appear to be more confused than frequency modulations.

To the best of our knowledge, this is the first completely unsupervised method for clustering similar signals. For modulation recognition, this method groups signal types by requiring only the signals themselves, without any *a priori* knowledge of the modulation classes.

Future plans include expanding the training dataset to 500,000 signals per modulation class, reducing the number of complex samples per signal from 4096 to 512, and using a ResNet-50 encoder adapted for time-domain data. We also plan to experiment with a DirectCLR architecture to prevent dimensional collapse during training [5].

Additionally, we aim to explore other architectures such as SWaV [6] for unsupervised model development on large radio datasets.

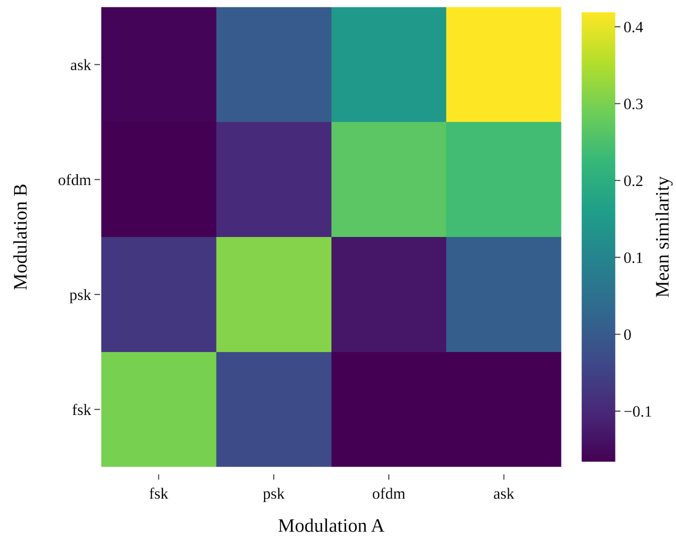

(a) Similarity matrix for grouped modulations including OFDM.

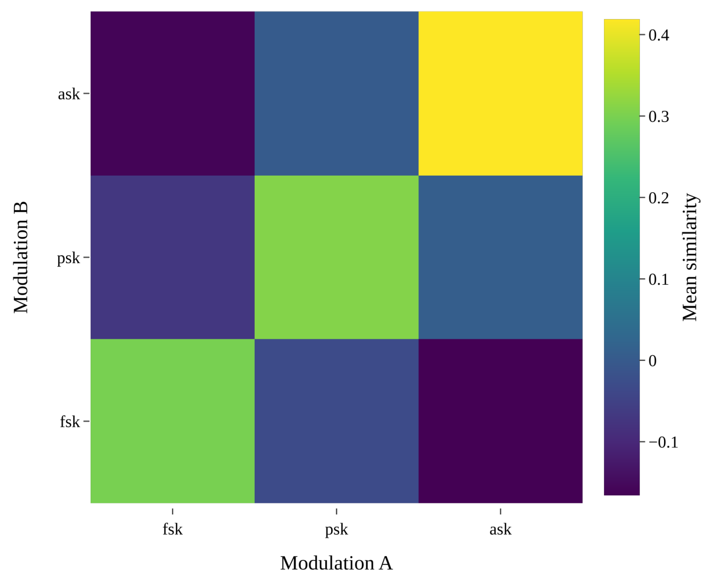

(b) Similarity matrix for grouped modulations excluding OFDM.

Figure 5.4: Similarity matrices for a simplified modulation recognizer formed by grouping modulations into parent classes.

References

[1] S. Harrison, R. Coles, T. Robishaw, and D. Del Rizzo, "RFI novelty detection using machine learning techniques," in *RFI2019: Coexisting with Radio Frequency Interference*,

2019.

[2] L. Boegner, M. Gulati, G. Vanhoy, P. Vallance, B. Comar, S. Kokalj-Filipovic, C. Lennon, and R. D. Miller, "Large scale radio frequency signal classification," 2022. arXiv:2207.09918.

[3] T. Chen, S. Kornblith, M. Norouzi, and G. Hinton, "A simple framework for contrastive learning of visual representations," June 2020. arXiv:2002.05709 [cs, stat].

[4] M. Tan and Q. V. Le, "EfficientNet: Rethinking model scaling for convolutional neural networks," *CoRR*, vol. abs/1905.11946, 2019.

[5] L. Jing, P. Vincent, Y. LeCun, and Y. Tian, "Understanding dimensional collapse in contrastive self-supervised learning," *arXiv preprint arXiv:2110.09348*, 2021.

[6] M. Caron, I. Misra, J. Mairal, P. Goyal, P. Bojanowski, and A. Joulin, "Unsupervised learning of visual features by contrasting cluster assignments," in *Proceedings of the 34th International Conference on Neural Information Processing Systems*, NIPS '20, (Red Hook, NY, USA), Curran Associates Inc., 2020.

Chapter 6

Conclusions

> You look at where you're going and where you are and it never makes sense, but then you look back at where you've been and a pattern seems to emerge.
>
> Robert M. Pirsig, Zen and the Art of Motorcycle Maintenance: An Inquiry Into Values

In this dissertation, I have presented the design, implementation, and evaluation of an end-to-end RFI monitoring system connecting modern, flexible RF front ends to the operational needs of radio observatories. Motivated by the increasing density and diversity of anthropogenic emitters and their threat to sensitive astronomical observations, the work has focused on building a practical monitoring chain: from wideband data acquisition and DoA estimation, through digital down-conversion and pooling architectures, to ML-based signal characterization and visualization. The individual contributions were developed as publishable components, but throughout the dissertation they are treated as interlocking pieces of a coherent monitoring pipeline as shown in fig. 6.1.

The first major technical contribution was the development of a wideband RFI monitoring instrument, with rigorously defined requirements and careful experimental assessment of receiver gain stability serving as its central elements. The commissioned instrument offers a calibrated sensitivity of $\approx -173\,\mathrm{dBm\,Hz^{-1}}$ from 350–$1800\,\mathrm{MHz}$. The requirements, design, and validation together constitute a practical template for future RFI monitor development.

Building on this foundation, the dissertation then addressed the problem of localizing RFI sources using an arbitrarily arranged array of antennas. I introduced a new DoA estimation method and showed that it consistently outperforms both subspace-based and image-domain algorithms.

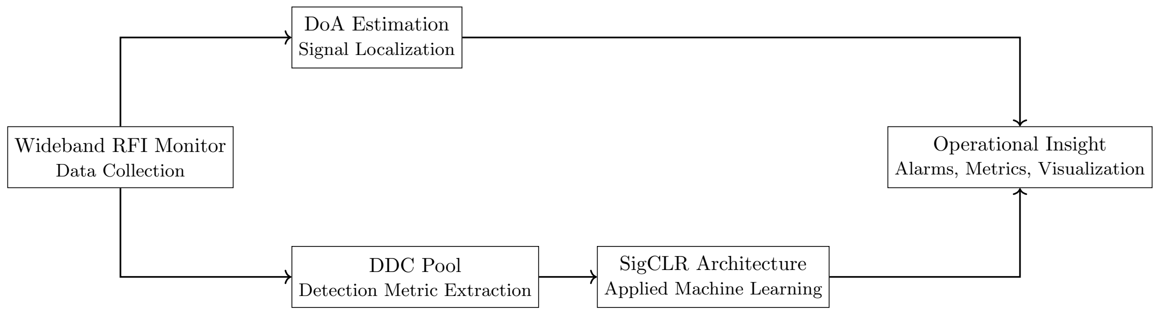

Figure 6.1: Block diagram of an end-to-end RFI monitoring system, which can continuously observe a wide frequency range, spatially localize signals, extract signal metrics from the data, and characterize the local RF environment.

The dissertation next focused on managing and processing the substantial data volumes generated by continuous wideband observations. To this end, I introduced a digital down-conversion architecture to flexibly channelize the spectrum and efficiently synthesize only the channels containing signals. Using this method, detected signals are synthesized into complex baseband time series from which signal metadata can be extracted, while the remainder of the detection-free spectrum can be discarded. This architecture enables long-duration monitoring campaigns without overwhelming storage or compute resources, and it provides a clean interface between the raw data output from the RFI monitor and higher-level classification algorithms.

The final technical element of the dissertation explored an unsupervised ML-based approach to signal characterization in the monitoring context. I introduced the SigCLR framework, a contrastive learning framework adapted from image-domain methods to learn robust representations of diverse signal types without requiring exhaustive manual labelling. While the framework was demonstrated on synthetic data, SigCLR is designed to accept data produced by the monitoring pipeline, supporting downstream tasks such as clustering and anomaly detection.

Viewed as a whole, the dissertation demonstrates that effective RFI monitoring is not a single algorithm or device, but a system-level effort that combines hardware design, signal processing, and machine learning. The wideband receiver provides direct, calibrated access to the spectrum; the DoA and DDC components transform those measurements into spatially localized, compactly represented signal records; and the SigCLR approach extracts higher-level structure from the resulting data. Together, these contributions constitute a monitoring pipeline that can provide continuous characterization of the RF environment and enable more informed responses to RFI events.

Beyond the concrete designs and algorithms, this work highlights several broader lessons for the community. First, modular architectures allow individual components—such as the ML-based characterization method—to be improved or replaced as technology advances, without requiring wholesale redesign. Second, the rapidly evolving RF landscape requires adaptable systems: systems capable of continuous learning—via online updates or periodic retraining—can more effectively accommodate new classes of interference.

In summary, this dissertation has proposed and validated an end-to-end RFI monitoring system combining modern RF hardware, digital signal processing, and machine learning. While much remains to be done, this is by design. A monitor of this kind must never be considered complete, but instead continually improved to better protect the scientific productivity of current and future radio observatories.

Chapter 7

Future Work

It never gets easier, you just go faster.

Greg LeMond

This chapter outlines both incremental improvements to the individual components introduced in Chapters 2 to 5 and broader efforts toward realizing a fully integrated, adaptive, and modular RFI monitoring architecture. Each area presents specific technical challenges and opportunities toward that goal.

7.1 Wideband RFI Monitor

The next steps for the wideband RFI monitor are to extend the architecture with an IF system to enable tuneable frequency coverage, and to implement an adaptive attenuation system that maintains linearity and dynamic range during strong RFI events, ensuring accurate signal characterization under varying interference conditions.

The length of coax between the antenna and the receiver should be reduced in order to improve the system temperature. Future iterations of the monitor should also digitize nearer to the receiver which will reduce the contribution to system temperature from the coax between the first and second stages of the signal chain.

Digitization nearer to the receiver will likely necessitate a different FPGA at which point an ADC with a higher effective number of bits (ENOB) and SFDR along with an internal mechanism to calibrate the effect of interleaving (which the current ADC lacks) should be considered.

7.2 DoA Estimation

Building multiple copies of the monitor hardware would enable spatial localization of signals of interest across the facility. Applying the DoA estimation method introduced in Chapter 3, which accommodates arbitrary array geometries, would allow antenna placement at the DRAO to be optimized subject to available power and network infrastructure constraints.

7.3 DDC Pool

Two clear future directions arise from Chapter 4. The first, discussed in detail in appendix 4.A, is to generalize the method to handle fractional oversampling ratios, which would reduce overall system data rates substantially. The second is to implement the DDC Pool as a real-time component of the DRAO RFI monitoring system, enabling direct interfacing with other pipeline components as signals are detected.

7.4 SigCLR Architecture

Chapter 5 introduced an initial concept and prototype for the SigCLR algorithm. To apply it effectively to real data, the training process must be further refined. So far, development attempts have used relatively large datasets, reducing the iteration rate. Future work will use smaller datasets and simpler model architectures to more rapidly converge on a robust solution for unsupervised modulation recognition and novelty detection. Once a satisfactory architecture is established, it must also be implemented for real-time operation.

7.5 Integration of *A Priori* Information

Integrating *a priori* information from external databases would enhance the utility of the monitoring pipeline. By cross-referencing real-time detections with known terrestrial and satellite allocations, the monitor can add metadata to each signal according to its regulatory compliance status. For instance, a signal detected via the DoA estimation module could be automatically correlated with the coordinates and frequency of a licensed local transmitter. Signals whose characteristics match the database parameters can be flagged as “authorized”, allowing the system to direct computational resources such as the DDC Pool toward unidentified or non-compliant signals.

7.6 Signal Detection

The work in this dissertation depends on signal detections being done on time-frequency data using an ML-network [1]. The network yields bounding boxes in time-frequency space that can then be passed to processing steps such as DoA estimation, a DDC Pool, modulation recognition, and a novelty detector. Improving the signal detection method is expected to improve generated signal metadata by providing more accurate bounding boxes to downstream tasks. Transformer and nearest-latent-neighbours architectures have shown promising results in this domain [2, 3] and will be explored. In order to validate the accuracy of these architectures, they will be compared with RFI excision data from on-site instruments such as CHIME. On-site experiments generate RFI masks using a variety of methods [4–8]. Comparing signal bounding boxes generated by the RFI monitor with excision masks generated by telescopes will be a valuable architecture validation method.

7.7 Tokenization of Raw Complex Voltage Streams

Conventional RFI mitigation typically operates on spectrograms or visibilities. As discussed in Chapter 4, the processes for generating these data inherently discard the time-domain structure that is highly effective at separating RFI from thermal noise. An avenue for future research is the direct tokenization of raw complex voltage streams. By utilizing a sliding-window, complex-valued tokenizer, the system can learn a discrete latent representation.

The operating principle is based on compressibility: thermal astronomical signals are Gaussian and incompressible, whereas anthropogenic signals have non-thermal structure. Consequently, any latent code that achieves compression on a data segment has successfully isolated RFI signals.

Although a standardized discrete-token scheme for raw RF voltage data does not yet exist, transferring representation strategies from adjacent domains represents an exciting frontier for unsupervised RFI mitigation. Crucially, the DDC Pool architecture developed in this dissertation is positioned to enable this idea. By efficiently extracting and preserving the isolated complex baseband waveforms of interference events, the DDC Pool provides the non-integrated, time-domain inputs required to train and deploy discrete-token models at scale.

7.8 Adapting Foundation Models

Recent years have seen unprecedented commercial investment in the development of generalized artificial intelligence (AI) systems. This rapid advancement has been driven by transformer architectures and tokenized modelling, culminating in highly capable large language models (LLMs). While industry has so far optimized for natural language, the underlying mechanics of these attention-based models have strong parallels with time-series RF data processing. Future iterations of RFI monitoring systems must adapt these commercially driven architectures. Transitioning from the bespoke learning frameworks presented in this work toward generalized foundation models offers a clear pathway to enhance the autonomous analysis and decision-making capabilities of RFI monitors.

7.9 Commercialization and Technology Transfer

The end-to-end RFI monitoring architecture presented in this work was deliberately designed with modularity and standardized data interfaces as foundational principles. While any monitor built according to this architecture must be commissioned specifically for the environment it will operate in, the hardware design, digital down-conversion pipeline, and machine learning frameworks are highly generalizable. A natural progression of this research involves the technology transfer and commercialization of these tools for a broader market. Packaging the system into a deployable, turnkey solution would allow other radio astronomy observatories, as well as national spectrum management authorities, to leverage autonomous, wideband RF characterization for regulatory enforcement and site protection.

7.10 Philosophy of RFI Monitoring Systems

The philosophy that an end-to-end RFI monitoring system should be modular, adaptive, and software-driven must guide all future development. This modularity must include external data streams; an intelligent monitor should understand the spectrum within the context of the broader regulatory environment. The archaic approach, which stored integrated spectra for later inspection, will prove increasingly inadequate as the RFI proliferates. The next generation of RFI monitoring systems must prioritize automation and continuous real-time data capture and processing.

The primary data products of an end-to-end RFI monitoring system will not be spectra or time series, but actionable outputs such as triggered alarms, collected metrics, and quantified interference statistics. By focusing on such outcomes, RFI engineers and scientists can build modules that transform spectral or time-series data into decision-ready information. Ultimately, this shift toward modular, real-time intelligence is the most promising path for protecting the research integrity of radio-quiet facilities. Implementing these strategies is a technical necessity, not merely an incremental improvement, for the continued development of terrestrial radio astronomy in an era of rapidly intensifying spectrum competition.

References

[1] S. Harrison, R. Coles, T. Robishaw, and D. Del Rizzo, "RFI novelty detection using machine learning techniques," in *RFI 2019*, 2019.

[2] M. Mesarcik, A.-J. Boonstra, E. Ranguelova, and R. V. van Nieuwpoort, "Learning to detect radio frequency interference in radio astronomy without seeing it," *Monthly Notices of the Royal Astronomical Society*, vol. 516, pp. 5367–5378, 11 2022.

[3] M. Mesarcik, "Hierarchical vision transformers for RFI mitigation in radio astronomy." RFI 2024 Conference.

[4] M. Rafiei-Ravandi and K. M. Smith, "Mitigating radio frequency interference in CHIME/FRB real-time intensity data," *The Astrophysical Journal Supplement Series*, vol. 265, p. 62, apr 2023.

[5] J. Taylor, N. Denman, K. Bandura, P. Berger, K. Masui, A. Renard, I. Tretyakov, and K. Vanderlinde, "Spectral kurtosis-based RFI mitigation for CHIME," *Journal of Astronomical Instrumentation*, vol. 08, no. 01, p. 1940004, 2019.

[6] L.-L. Zhao, X.-H. Luan, X. Chao, Y.-C. Wang, J.-K. Li, Z.-Z. Tao, T.-J. Zhang, H.-F. Wang, and D. Werthimer, "An improved machine learning approach for radio frequency interference mitigation in FAST–SETI survey archival data," *The Astronomical Journal*, vol. 171, p. 86, jan 2026.

[7] J. Bai, S. Dai, N. Wang, S. Osłowski, S. Wang, G. Hobbs, J. Yuan, W. Yan, Q. Zhi, L. Shang, and et al., "Radio frequency interference identification using eigenvalue

decomposition for multi-beam observations," *Publications of the Astronomical Society of Australia*, vol. 43, p. e008, 2026.

[8] R. S. Hombal, L. Levin, B. W. Stappers, M. Droog, A. Karastergiou, D. Lumbaa, M. B. Mickaliger, A. Naidu, K. M. Rajwade, J. Sepulveda, B. Shaw, S. Singh, and T. Prabu, "Evaluating the effectiveness of radio frequency interference removal algorithms for single pulse searches," *RAS Techniques and Instruments*, vol. 5, p. rzag004, 01 2026.