Design and Development of an Instructional Interferometer for Radio Observations

Thesis

Submitted in fulfillment of the requirements
for the award of the degree of

Master of Science

by

Sangeetha A



Supervisor

Prof. Abhirup Datta

Department of Astronomy, Astrophysics and Space Engineering

Indian Institute of Technology Indore

May 2026

INDIAN INSTITUTE OF TECHNOLOGY INDORE

CANDIDATE'S DECLARATION

I hereby certify that the work which is being presented in the thesis entitled **Design and Development of an Instructional Interferometer for Radio Observations** in the partial fulfillment of the requirements for the award of the degree of **MASTER OF SCIENCE** and submitted in the **Department of Astronomy, Astrophysics and Space Engineering, Indian Institute of Technology Indore**, is an authentic record of my own work carried out during the time period from July 2024 to May 2026 under the supervision Prof. Abhirup Datta, Professor, Department of Astronomy, Astrophysics and Space Engineering, IIT Indore.

The matter presented in this thesis has not been submitted by me for the award of any other degree of this or any other institute.

A. Datta
07/05/2026

Signature of the student with date
(Sangeetha A)

This is to certify that the above statement made by the candidate is correct to the best of my/our knowledge.

07/05/2026

Signature of the Supervisor of
M.Sc. thesis (with date)
(Prof. Abhirup Datta)

SANGEEETH A has successfully given his/her M.Sc. Oral Examination held on **04-05-2026**.

Signature(s) of Supervisor(s) of MSc thesis
Date: 07/05/2026

Manoneeta Chakraborty
Programme Coordinator, M.Sc.
Date: 08/05/2026

Manoneeta Chakraborty
Convenor, DPGC
Date: 08/05/2026

S. Das
HoD, DAASE
Date: 08/05/26

Abstract

Radio interferometry is central to modern high-resolution radio astronomy. However, practical experience with this technique remains limited in many educational settings because the required hardware is often proprietary and the associated signal chains are poorly documented. To address this gap, a two-element digitally correlating interferometer was constructed and evaluated at 1.420 GHz. The system incorporated matched pyramidal horn antennas (14.2 dBi gain), a front-end chain optimized using the Friis noise formula (approximately 68 dB total gain), and a USRP B210 software-defined radio backend. An FX correlator, implemented in GNU Radio with Python UHD scripting, enabled real-time spectral cross-correlation while preserving the full complex visibility. Phase coherence was validated in controlled bench tests using injected tones across a four-element, six-baseline configuration. Two-element solar drift-scan observations at four east-west baselines consistently detected interferometric fringes. The measured fringe periods were consistent with Earth-rotation predictions within 16%, and the fringe count increased monotonically with baseline length, ranging from 3 to 9 peaks over a fixed 3-hour window. The visibility amplitude remained broadly constant across sessions, indicating that the Sun is unresolved at all baselines used (null baseline 29.5 m). All stages of the system, from antenna design and inter-device synchronisation to correlator implementation, are documented in detail to support reproducibility and to facilitate future use by others.

Acknowledgements

I would like to express my sincere gratitude to my supervisor, Prof. Abhirup Datta, for his guidance, patience, and encouragement throughout this project. His insightful discussions and unwavering support enabled me to overcome numerous challenges.

I am also grateful to Dr. Harsha Avinash Tanti for his valuable suggestions and his readiness to assist with technical difficulties related to radio astronomy instrumentation. I extend my thanks to our lab assistant, Mr. Abhijeet Dutta, for his assistance in arranging facilities and resources.

I appreciate the support and helpful discussions from my seniors, fellow students, and labmates, especially Sudhamshu G, Jibin Jose, Kaivan Sanjay Shah, Nityananda Padhi and Anismita Biswas, whose camaraderie made the long hours in the lab more enjoyable. I also thank my friends from M.Tech. for their assistance in setting up equipment for rooftop observations.

I am grateful to the Department of Astronomy, Astrophysics and Space Engineering at IIT Indore for providing the laboratory facilities, equipment, and computational resources essential for this project.

Lastly, I would like to express my gratitude to my family, brothers, and friends for their unwavering encouragement and support throughout this academic journey.

Contents

Abstract i

Acknowledgements ii

1 Introduction 1

1.1 The Radio Universe: Origins and Insights . . . 1

1.1.1 Unique Capabilities of Radio Astronomy . . . 2

1.2 From Single Dishes to Interferometry . . . 3

1.2.1 The Resolution Problem . . . 3

1.2.2 The Interferometry Solution . . . 4

1.2.3 Modern Interferometric Arrays . . . 4

1.3 The Gap This Project Fills . . . 5

1.3.1 Limitations of Existing Educational Interferometers . . . 6

1.3.2 Motivation and Novelty . . . 6

1.4 Research Objectives . . . 7

1.5 Thesis Organisation . . . 7

2 Foundations of Radio Interferometry 8

2.1 Single-Dish Radio Astronomy . . . 8

2.1.1 Collecting Radio Radiation . . . 8

2.1.2 Angular Resolution . . . 9

2.1.3 The Radiometer Equation . . . 9

2.1.4 Limitations of Single Dish . . . 10

2.2 Coordinate Systems . . . 11

2.2.1 Equatorial Coordinates: Right Ascension and Declination . . . 11

2.2.2 Horizontal Coordinates and Converting Between Systems . . . 11

2.3 The Two-Element Interferometer . . . 12

2.3.1 Geometric Delay . . . 12

2.3.2 The Visibility Function . . . 13

2.3.3 The van Cittert–Zernike Theorem: Connecting Fringes to Sky . . . 13

2.4 The UV Plane and Aperture Synthesis . . . 14
2.4.1 The UV Plane . . . 14
2.4.2 Aperture Synthesis and Earth-Rotation . . . 14
2.5 The Fringe Pattern . . . 14
2.5.1 Fringe Period . . . 15
2.5.2 Fringe Amplitude and Source Structure . . . 15
2.6 Mathematical Tools: Fourier Transform, Correlation, and Convolution . . . 16
2.6.1 The Fourier Transform . . . 16
2.6.2 Cross-Correlation . . . 17
2.6.3 Convolution . . . 18

3 Hardware Design: Antenna and Front-End Electronics 20
3.1 System Overview . . . 20
3.2 Antenna Design and Characterisation . . . 21
3.2.1 Design Goals . . . 22
3.2.2 Analytical Design Equations . . . 22
3.2.3 10 dBi Compact Design . . . 23
3.2.4 Horn Antenna: Dimensions and Predicted Performance . . . 24
3.2.5 Antenna Simulation in CST Studio Suite . . . 25
3.3 Front-End Electronics . . . 28
3.3.1 Components . . . 29
3.3.2 Noise Figure and the Friis Formula . . . 32
3.3.3 Front-End Configurations Evaluated . . . 33
3.3.4 Front-End Configuration Selection . . . 35

4 Digital Backend 37
4.1 Backend 1: RTL-SDR . . . 37
4.1.1 Motivation . . . 37
4.1.2 The Phase Coherence Requirement . . . 38
4.1.3 Bench Test Setup . . . 39
4.1.4 GNU Radio Flowgraph . . . 39
4.1.5 Result and Interpretation . . . 40
4.2 Backend 2: USRP B210 Software-Defined Radio . . . 40
4.2.1 Hardware Overview . . . 40
4.2.2 Single USRP B210: Two-Element Phase Coherence Verification . . . 42
4.2.3 Extending to Four Elements: Two Synchronised USRP B210s . . . 44
4.2.4 Synchronisation: Issues Encountered and Solutions . . . 48
4.2.5 FFT Flowgraph: Baseline 1–3 (Two-USRP, Single Baseline) . . . 50

4.2.6 Four-Element Six-Baseline Configuration . . . 51
4.2.7 Current Status and Limitations of the Four-Element System . . . 54
4.3 Outdoor Solar Drift Scan Observations . . . 54
4.3.1 Observation Setup . . . 55
4.3.2 Front-End Signal Chain for Outdoor Observations . . . 56
4.3.3 Digital Back-End: FX Correlator . . . 57

5 Results **59**
5.1 RTL-SDR Backend: Phase Coherence Test . . . 59
5.1.1 Cross-Correlation Result . . . 59
5.1.2 Limitations of Daisy-Chain . . . 60
5.2 USRP B210 Backend: Phase Synchronisation Results . . . 60
5.2.1 Stage 1: Single USRP B210 — Internal Phase Coherence . . . 60
5.2.2 Stage 2: Two-USRP Inter-Device Synchronisation . . . 61
5.2.3 Physical Interpretation of the Phase Behaviour . . . 63
5.2.4 Stage 3: Four-Element Six-Baseline Validation . . . 64
5.2.5 USRP Synchronisation Summary . . . 67
5.2.6 Representative Solar Interferometric Fringe Observations . . . 67
5.2.7 Baseline $B = 1.00$ m ($u = 4.74\,\lambda$) . . . 68
5.2.8 Baseline $B = 1.88$ m ($u = 8.9\,\lambda$) . . . 68
5.2.9 Baseline $B = 2.00$ m ($u = 9.5\,\lambda$) . . . 69
5.2.10 Baseline $B = 2.40$ m ($u = 11.4\,\lambda$) . . . 70
5.3 Dependence of Fringe Properties on Baseline . . . 71

6 Summary and Future Scope **74**
6.1 Summary of the Project . . . 74
6.1.1 Instrument Design and System Architecture . . . 74
6.1.2 Experimental Validation and Observational Results . . . 74
6.1.3 Key Lessons . . . 75
6.1.4 Limitations . . . 75
6.2 Future Scope . . . 76
6.2.1 On-Sky Validation of the Four-Element System . . . 76
6.2.2 RFSoC-Based Backend . . . 76
6.2.3 Additional Astronomical Targets . . . 77
6.2.4 Aperture Synthesis . . . 77

References **79**

List of Figures

1.1 Karl Jansky with his antenna that first detected cosmic radio waves in 1932. Built on wheels salvaged from a Model T Ford, this rotating structure discovered radio emission from the centre of the Milky Way. Image credit: National Radio Astronomy Observatory (NRAO) [2]. 1

1.2 Grote Reber's 9.6-metre parabolic radio telescope in Wheaton, Illinois (1937), widely regarded as the world's first dedicated radio astronomy instrument. Reber's surveys produced the first radio continuum maps of the Milky Way. Image credit: National Radio Astronomy Observatory (NRAO). [4]. 2

1.3 The 100-metre Robert C. Byrd Green Bank Telescope (GBT), the world's largest fully steerable single-dish radio telescope. At a wavelength of 21 cm, its angular resolution ($\sim \lambda/D$) is insufficient to resolve fine details in most cosmic sources, motivating the use of interferometry. Image credit: NSF/Green Bank Observatory (GBO) [5]. 4

1.4 The Very Large Array (VLA) in New Mexico, with 27 antennas working together as an interferometer and achieving the angular resolution of a 36-kilometre telescope. Image credit: NRAO [7]. 5

2.1 Radiation pattern of a directional antenna showing the main lobe, sidelobes, and back lobe. The angular width of the main lobe defines the effective beamwidth of the antenna, within which incoming radiation is integrated. The half-power beamwidth (HPBW) corresponds to the angular width of the main lobe at half maximum power. Image adapted from Wikimedia Commons [11]. 10

2.2 The equatorial (RA, Dec) and horizontal (Alt, Az) coordinate systems. As the Earth rotates, a source's RA and Dec remain fixed, while its altitude and azimuth change continuously 12

3.1 Block diagram of the n-element instructional radio interferometer. Each antenna feeds an identical front-end module and a synchronised SDR/RFSoC backend. All backends connect to a central laptop for software correlation and fringe analysis. The dashed box highlights the clock-synchronisation requirement that motivated the choice of digital backend, as discussed in Chapter 4. 21
3.2 Flare section of the horn, showing the aperture width W_a, aperture height H_a, and flare length L. Together with the throat dimensions, these parameters fully define the antenna geometry. 26
3.3 Waveguide throat and feed transition, showing the broad dimension a_{th}, narrow dimension b_{th}, and the feed point where the coaxial connector excites the TE_{10} mode. 26
3.4 Simulation model of the laboratory horn, illustrating multiple views of the geometry, including the aperture, flare section, and feed region. 27
3.5 Simulated return loss S_{11} of the laboratory horn antenna. The minimum value of -9.22 dB at 1.42 GHz indicates good impedance matching near the design frequency. 27
3.6 Simulated 3D gain pattern at 1.42 GHz. The peak gain of $\approx$ 14.2dBi agrees with the analytical prediction of Equation 3.10, validating the use of $\eta = 0.65$. 28
3.7 Simulated far-field radiation pattern (|Directivity|) at 1.42 GHz for the $\phi = 90°$ plane. The main lobe is centered at 0° with a peak directivity of 14.6dBi. The half-power beamwidth (3dB) is approximately 35.3°, and the side lobe level is -21.4 dB, consistent with the expected directional characteristics of the antenna. 28
3.8 The Nooelec SAWbird+ H1 LNA. The module provides a gain of approximately 39.5 dB while contributing very little additional noise (NF $\approx$ 0.04 dB). Its integrated SAW pre-filter also helps suppress strong out-of-band interference before amplification. 29
3.9 Measured gain and noise figure of the SAWbird+ H1 as functions of frequency. Both peak near 1.42 GHz, confirming that the module is well suited to hydrogen-line observations. 30
3.10 Mini-Circuits VBFZ-1400-S+ band-pass filter, which isolates the 1350–1450 MHz hydrogen-line window. 30
3.11 S-parameter response of the VBFZ-1400-S+. The 2dB in-band insertion loss and >40dB out-of-band rejection are clearly visible. Because this filter is placed after the first LNA, its 2dB loss contributes negligibly to the overall system noise figure (Section 3.3.2). 31
3.12 Mini-Circuits SHP-1000+ high-pass filter, placed at the end of the chain to block FM radio and other sub-GHz interference. 31

3.13 Measured response of the SHP-1000+. Strong attenuation below 1 GHz and low insertion loss in the 1–2 GHz band are both visible. 32
3.14 VNA measurement of Configuration 1: LNA → BPF → LNA → HPF. The LNA-first ordering gives the lowest system noise figure (Friis formula, Eq. 3.15) and was selected for our system. 34
3.15 VNA measurement of Configuration 2: BPF → LNA → LNA → HPF. Placing the filter first protects against strong RFI but raises the noise floor by ≈ 2 dB. . 34
3.16 VNA measurement of Configuration 3: LNA → BPF → Attn → LNA → HPF. The attenuator improves stability for long observations. System noise figure impact is negligible because the attenuator sits after the first LNA. 35
3.17 VNA-measured gain comparison of the three front-end configurations. 36

4.1 Clock routing on the RTL-SDR Blog V3. A shared reference frequency can be provided in hardware, but as the bench tests show, this does not guarantee phase coherence [15]. 38
4.2 Bench test setup for the RTL-SDR phase stability measurement. Two RTL-SDR dongles receive the same split signal via a two-way RF power splitter. Both are connected to the host computer through a powered USB hub. Any variation in the measured phase difference between the two channels is attributable entirely to the receivers. 39
4.3 GNU Radio flowgraph for the RTL-SDR phase coherence test. Both dongles stream into separate FFT paths, after which a Multiply Conjugate block forms the cross-power spectrum. The magnitude and phase of the cross-product are displayed in real time. 40
4.4 The Ettus Research USRP B210. Both receive channels share a single AD9361 transceiver chip and internal clock, providing inherent two-channel phase coherence. The REF IN and PPS IN SMA connectors allow locking to an external frequency reference and pulse-per-second timing signal for multi-device synchronisation. 41
4.5 Single USRP B210 phase coherence test: (a) setup overview; (b) RX2 input connections. 43
4.6 GNU Radio Companion FX correlator flowgraph for single USRP B210 phase coherence verification. Both channels are processed through parallel FFT chains, and the cross-power spectrum is computed using a Multiply Conjugate block. The resulting visibility magnitude and phase are displayed in real time. 44

4.7 GNU Radio FX correlator flowgraph for the two-USRP incoherence baseline test. Both USRP sources operate on their internal clocks with no synchronisation (Clock Source: Internal, Sync: No Sync). Each channel passes through an 8192-point Blackman–Harris FFT before cross-multiplication. The cross-power magnitude and phase are displayed in real time. 45
4.8 Moku:Pro waveform generator configuration used for synchronisation. The upper panel shows the 10 MHz reference signal, while the lower panel shows the 1 Hz PPS signal with a DC offset producing a 0–3.3 V waveform. Both outputs share a common internal timebase. 46
4.9 GNU Radio flowgraph for two-USRP time-domain synchronisation testing. Each USRP is accessed through a separate UHD source block with external clocking and PPS. The signals are processed in parallel, and their real, imaginary, magnitude, and phase components are displayed and recorded. 47
4.10 Terminal output confirming both USRP B210 units are operating over USB 3.0. 49
4.11 Bench setup overview for the two-USRP synchronisation test. (a) Overall system layout. (b) RF signal injection path. 50
4.12 Rear panel of one USRP B210 showing the USB 3.0 data connection, the 10 MHz REF IN input, and the PPS IN input. The REF IN and adjacent GPS port are physically identical SMA connectors; connecting to the wrong port is a common source of error during synchronisation setup. 50
4.13 GNU Radio Companion flowgraph for frequency-domain synchronisation verification of Baseline 1–3. Both USRP B210 units stream Channel 0 with external clock and time sources. Signal flow: Stream to Vector → FFT (8192-point, Blackman-Harris) → Vector to Stream → Multiply Conjugate → Complex to Mag Phase → QT GUI Time Sink. The cross-correlation output from this flowgraph is presented in Chapter 5 51
4.14 Overall bench setup for the four-element interferometer test, showing the USRP B210 units, signal generator, RF 4-way splitter, and reference timing distribution via the Moku:Pro. 52
4.15 Close-up of the four USRP receive ports used in the four-element test. The RF A RX2 and RF B RX2 inputs on both USRP 1 and USRP 2 are each connected to one output of the four-way RF splitter, ensuring identical signal injection across all channels. 53

4.16 GNU Radio four-element FX correlator flowgraph. Four independent FFT paths (two channels from each USRP) feed into six Multiply Conjugate blocks computing all six baselines simultaneously. The magnitude and phase of all six baselines are displayed in real time and written to file sinks. The results from this flowgraph are presented in Chapter 5. 54
4.17 Comparison of observing configurations. Moving to the rooftop improved data quality and reduced obstruction effects. 56
4.18 Measured gain (left) and noise figure (right) for the two front-end chains. The shaded region marks the H I band around 1.42 GHz. Both chains show closely matched responses within the observation band. 57
4.19 GNU Radio FX correlator pipeline. The two input channels are Fourier transformed and cross-correlated to obtain complex visibility. 57

5.1 Cross-correlation result for two daisy-chained RTL-SDR dongles receiving the same split tone. The phase (upper) spans $[-\pi, +\pi]$ with no stable structure, while the amplitude (lower) remains near zero except for a brief startup transient around $t \approx 2800$–$3400\,\mu s$. No stable cross-correlation phase is established, showing that the daisy-chain provides a common reference frequency but not a shared phase. 59
5.2 Single USRP B210 internal coherence test. The upper panel shows the instantaneous cross-correlation phase, while the lower panel shows the corresponding amplitude. A strong correlated signal appears as a distinct peak in amplitude. The phase fluctuates because of thermal noise and finite signal-to-noise ratio, but remains statistically stable over time. 61
5.3 Two-USRP cross-spectrum without external synchronisation. The cross-correlation phase (upper panel) is distributed across $[-\pi, +\pi]$ with no stable structure, indicating complete phase incoherence. The amplitude (lower panel) shows only weak and inconsistent correlation, confirming that the two devices are not phase-locked. 62
5.4 Two-USRP cross-correlation with full synchronisation. The phase (upper panel) remains nearly constant at ~ -2.8 rad with only small fluctuations, indicating successful phase locking. The magnitude (lower panel) is stable and non-zero, confirming continuous correlated signal detection. 62

5.5 FFT cross-spectrum with synchronisation enabled. The phase (upper panel) remains bounded with no systematic drift, while the amplitude (lower panel) shows a clear correlated peak at the injected tone frequency. The noise bins remain uncorrelated, confirming that the observed structure arises from the signal rather than from system artefacts. 63
5.6 Four-element, six-baseline cross-correlation without external synchronisation. The phase (top) is erratic across all baselines, with no stable structure, while the amplitude (bottom) remains coherent only for the internal Baseline 3–4 within USRP 2. All cross-device baselines are incoherent, confirming the need for external synchronisation. 64
5.7 Overview of all six baselines with synchronisation active. Phase (upper) and amplitude (lower) versus time. Colour coding: 1–2 (blue), 1–3 (red), 1–4 (green), 2–3 (black), 2–4 (cyan), 3–4 (magenta). All phase traces remain bounded and stable, and correlated amplitude spikes appear simultaneously across all baselines near $t \approx 3000$ µs. 65
5.8 Individual cross-correlation phase (upper panel) and amplitude (lower panel) for all six baselines. The two internal baselines (1–2 and 3–4) show the cleanest phase behaviour, as both channel pairs share a single AD9361 chip. All four cross-device baselines (1–3, 1–4, 2–3, 2–4) show bounded phase with no sweeping and correlated amplitude spikes appearing in the same ∼2800–3400 µs window, confirming phase-coherent operation across every inter-device channel combination. 66
5.9 Solar drift-scan observation at baseline $B = 1.00$ m ($u = 4.74\,\lambda$). The upper panel shows the cross-power spectrum waterfall over the 3-hour session. The middle panel shows the visibility magnitude as a function of time, with three fringe peaks clearly identified (marked by triangles); the measured period is 44.4 ± 2.8 min. The lower panel shows the corresponding unwrapped visibility phase. 68
5.10 Solar drift-scan observation at baseline $B = 1.88$ m ($u = 8.9\,\lambda$). Six fringe peaks are identified over the 3-hour interval, giving a measured period of 25.6 ± 1.5 min, in agreement with the session-averaged prediction of 26.5 min to within 3.4%. . 69
5.11 Solar drift-scan observation at baseline $B = 2.00$ m ($u = 9.5\,\lambda$). Six fringe peaks are identified over the 3-hour session, giving a measured period of 23.2 ± 1.8 min compared with the session-averaged prediction of 24.9 min (agreement within 7.3%). This session gave the highest median visibility magnitude (31.5 dB) and the lowest phase scatter (0.087 rad) of all four sessions. 70

5.12 Solar drift-scan observation at $B = 2.40$ m ($u = 11.4\ \lambda$). Nine fringe peaks are detected over 3 hours, consistent with the expected count of 8.7. The measured period, 17.5 ± 2.0 min, is $\sim 16\%$ shorter than the predicted 20.8 min, likely due to SNR-related peak selection bias and increased baseline uncertainty at longer separations. 71

6.1 AMD Xilinx ZCU216 RFSoC development board used for initial familiarisation and spectrometer development. Preliminary laboratory tests included board bring-up, ADC verification, and integration with the CASPER toolflow. 77

List of Tables

3.1 Comparison of the compact 10 dBi design target with the horn antenna used in all observations. Both are designed to operate at $f_0 = 1.420$ GHz. 25

4.1 Key specifications of the USRP B210 relevant to this project. 42
4.2 Channel mapping and baseline definitions for the four-element system. Internal baselines are inherently phase-coherent; cross-device baselines require external synchronisation. 52

5.1 Summary of fringe parameters for the four representative observation sessions. Each session had a duration of 3 hours at azimuth 180° and elevation $\approx 70°$. Predicted periods use the session-averaged hour angle correction $\langle \cos H \rangle = 0.974$ (Eq. 5.1). Agreement is defined as $|T_{\text{meas}} - \langle T_{\text{pred}} \rangle| / \langle T_{\text{pred}} \rangle \times 100\%$. 72

Chapter 1

Introduction

1.1 The Radio Universe: Origins and Insights

Radio astronomy did not originate with astronomers, but with an engineer named Karl Jansky. In 1932, while studying the static that disrupted transatlantic telephone transmissions, he detected a faint, constant hiss that appeared to drift slowly across the sky each day. After months of systematic measurements, he concluded that the signal did not come from Earth or the Sun, but instead from the centre of our own Milky Way galaxy [1]. At the time, most professional astronomers were largely unimpressed. They assumed that stars were far too cool to produce radio emission strong enough to be detected.

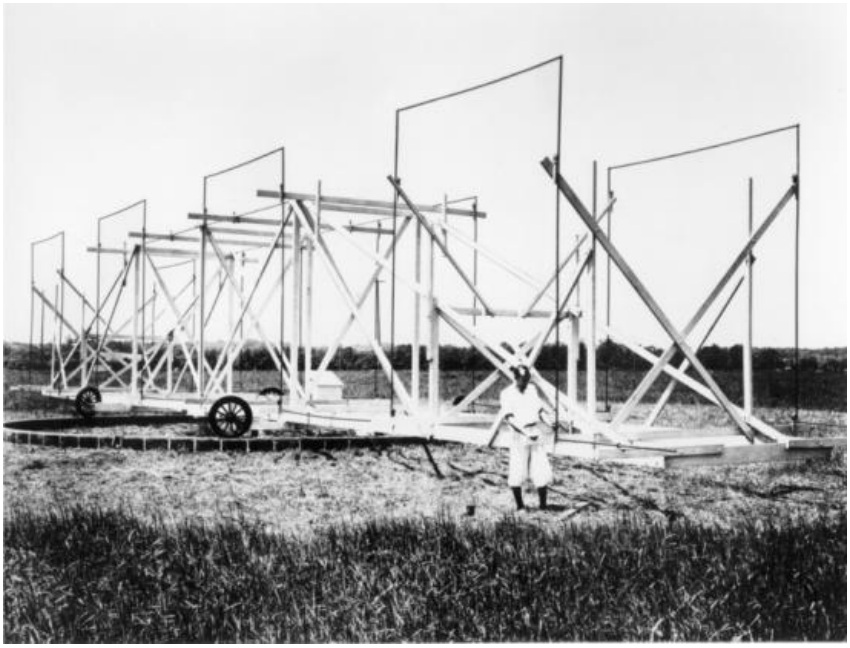

Figure 1.1: Karl Jansky with his antenna that first detected cosmic radio waves in 1932. Built on wheels salvaged from a Model T Ford, this rotating structure discovered radio emission from the centre of the Milky Way. Image credit: National Radio Astronomy Observatory (NRAO) [2].

The field might have stayed dormant had it not been for Grote Reber, a radio engineer who understood the significance of what Jansky had found. In 1937, Reber built the world's first parabolic radio telescope entirely at his own expense, in his backyard in Wheaton, Illinois. His 9.6-metre dish offered far better sensitivity and directionality than Jansky's rotating array. He initially observed at 3300 MHz, expecting thermal emission from stars, and detected nothing. Undeterred, he dropped his frequency to 160 MHz and finally succeeded in 1938. Between 1938 and 1943, Reber produced the first comprehensive radio survey of the sky, mapping the Galactic plane and identifying bright discrete sources including Cas A and Cyg A. These results, published in the *Astrophysical Journal*, helped establish radio astronomy as a recognised scientific field and provided some of the first clear evidence for non-thermal emission from cosmic sources [3].

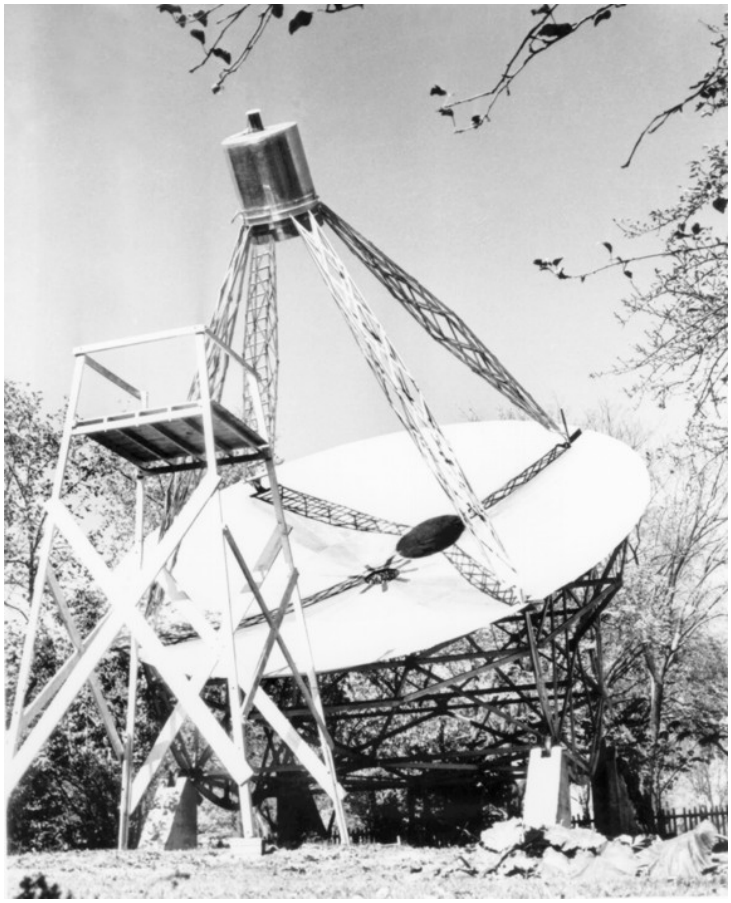

Figure 1.2: Grote Reber's 9.6-metre parabolic radio telescope in Wheaton, Illinois (1937), widely regarded as the world's first dedicated radio astronomy instrument. Reber's surveys produced the first radio continuum maps of the Milky Way. Image credit: National Radio Astronomy Observatory (NRAO). [4].

1.1.1 Unique Capabilities of Radio Astronomy

Radio observations reveal phenomena that are entirely inaccessible to optical telescopes. One clear example is the 1.420 GHz hydrogen line used in this project, where neutral hydrogen emits

at this frequency through the hyperfine spin-flip transition and has no optical counterpart. Cold hydrogen clouds at temperatures of 10–100 K are therefore invisible to optical instruments, yet they produce detectable emission at radio wavelengths. As a result, radio observations provide the only practical means of mapping the large-scale distribution of neutral gas in the Galaxy [1].

A second advantage that motivated the choice of 1.42 GHz for this instrument is the relative spectral cleanliness near the hydrogen line. International regulations protect this band from transmissions, thereby reducing local interference compared with adjacent frequencies. In practice, some residual interference from mobile networks and radar was still present in our observations (Section 3.3), but the protected status of the band made the signal chain substantially easier to design than it would have been at an unprotected frequency.

More broadly, radio wavelengths penetrate dust that obscures optical light, enabling observations of regions such as galactic centres and star-forming nebulae. Radio observations can also be conducted in daylight, which was essential for the midday solar measurements in this project. The radio spectrum spans approximately 3 kHz to 300 GHz and is sometimes extended into the submillimetre regime up to approximately 1 THz, providing access to phenomena ranging from synchrotron emission and free–free emission to the Cosmic Microwave Background [1]. The 1.42 GHz band is particularly suitable for this work because it is accessible with low-cost software-defined radio hardware.

1.2 From Single Dishes to Interferometry

1.2.1 The Resolution Problem

Every telescope has a fundamental limit on the finest detail it can resolve, known as its angular resolution. For a single dish, this is described by the simple relation:

$$\theta \approx \frac{\lambda}{D} \tag{1.1}$$

where θ is the smallest angular separation the telescope can distinguish, λ is the observing wavelength, and D is the dish diameter. This equation shows a major practical difficulty. At a wavelength of 21cm, corresponding to the neutral hydrogen line, achieving an angular resolution of 1arcsecond, comparable to a modest optical telescope, would require a dish nearly 50kilometres in diameter. Even the 100-metre Green Bank Telescope, the world's largest fully steerable single dish, can resolve only features larger than about 7arcminutes at this wavelength, which is far too coarse for studying fine structure in most cosmic sources.



Figure 1.3: The 100-metre Robert C. Byrd Green Bank Telescope (GBT), the world's largest fully steerable single-dish radio telescope. At a wavelength of 21 cm, its angular resolution ($\sim \lambda/D$) is insufficient to resolve fine details in most cosmic sources, motivating the use of interferometry. Image credit: NSF/Green Bank Observatory (GBO) [5].

1.2.2 The Interferometry Solution

Radio astronomers addressed this problem with an elegant solution. Instead of building one extremely large dish, they combine multiple smaller antennas spread across a wide area and process their signals together. The resolution of this combined system is determined not by the size of each individual antenna, but by the largest separation between them, called the *baseline*. Two antennas separated by one kilometre can therefore achieve the angular resolution of a one-kilometre dish. When signals from two antennas are brought together and multiplied in a process called *correlation*, the output is a slowly varying oscillatory pattern known as a *fringe*. This fringe arises because the same wavefront from a cosmic source reaches the two antennas at slightly different times, depending on the source position in the sky. As the Earth rotates, this delay, and the associated phase difference changes continuously and produce the observed oscillation. The fringe rate and amplitude, therefore, contain information about the source position and its angular structure [6]. By combining measurements from many baselines, a detailed image of the source can be reconstructed using aperture synthesis.

1.2.3 Modern Interferometric Arrays

This principle underlies all major radio observatories in use today. The Very Large Array (VLA) in New Mexico consists of 27 antennas arranged in a Y-shaped configuration, with a maximum baseline of 36 kilometres, giving it a resolution thousands of times finer than that of any single dish. The Giant Metrewave Radio Telescope (GMRT) near Pune, India, uses 30 antennas in a

hybrid arrangement that is sensitive to both large-scale structure and fine detail. The Atacama Large Millimetre/submillimeter Array (ALMA) in Chile operates 66 high-precision antennas to study the cold molecular universe. These instruments have transformed our understanding of phenomena ranging from planet-forming disks to the earliest galaxies in the universe.



Figure 1.4: The Very Large Array (VLA) in New Mexico, with 27 antennas working together as an interferometer and achieving the angular resolution of a 36-kilometre telescope. Image credit: NRAO [7].

The principles behind these instruments are well established and deeply important. However, for most students and early researchers, they often remain abstract, something to read about rather than experience directly.

1.3 The Gap This Project Fills

A single radio dish is analogous to a human ear: it can detect a signal and measure its strength, but cannot determine fine spatial detail about the source. An interferometer is more like two ears separated by a large distance. By comparing the signals received at each antenna and the time difference between them, finer angular structure can be resolved than either antenna could detect alone. The greater the separation, the finer the resolvable detail.

Most educational radio astronomy projects use single-dish instruments. The Small Radio Telescope uses a satellite dish for basic sky surveys [8], and horn antenna systems combined with software-defined radios have been used to detect the 21 cm hydrogen line and measure Galactic rotation curves [9]. These are useful for learning the basics of radio detection, but they stop short of the step that underlies all modern radio astronomy: combining signals from two or more antennas to measure interferometric visibilities.

For many students, the signal path between an antenna and a visibility measurement remains opaque. The correlation step, which is central to interferometry, is typically embedded in expensive specialised hardware that can be operated but not easily inspected or modified.

1.3.1 Limitations of Existing Educational Interferometers

A small number of educational interferometers have been described in the literature, but each has limitations that this project aims to overcome.

Koda et al. built a Michelson-type interferometer at 11 GHz using commercial satellite hardware and demonstrated resolved observations of the Sun [10]. However, the system depends on proprietary components whose internal operation is not transparent, which reduces its value as a teaching instrument.

Pandian et al. at Harvard developed a two-element adding interferometer using horn antennas and SDR receivers, and successfully detected solar fringes at 1.42 GHz [9]. Because an adding interferometer combines signals before digitisation, it does not provide direct access to calibrated complex visibility phase.

Several SDR-based FX correlators have also been demonstrated in laboratory settings, but these generally do not extend to a complete end-to-end system that includes antenna design, calibrated observations, and on-sky validation.

None of these instruments provides a fully open, digitally correlating interferometer in which every stage, from antenna design to correlator implementation, is explicitly documented and reproducible with modest laboratory resources.

1.3.2 Motivation and Novelty

The primary motivation for this work was the lack of scalable and reproducible educational interferometer platforms that can provide an end-to-end understanding of radio instrumentation, synchronisation, and interferometric signal processing. Although radio interferometry is often introduced in largely theoretical terms, many of its practical aspects become much clearer through direct experimental implementation. This project therefore focused on the development of a compact SDR-based multi-element radio interferometer operating at 1.42 GHz. The system encompasses antenna design, RF front-end characterisation, backend integration, synchronisation, and FX correlation implemented in GNU Radio. By operating directly on digitised IQ samples, the correlator preserves the full complex visibility, including both amplitude and phase information. A central emphasis of the project was reproducibility. Key design choices, including antenna geometry, front-end architecture, synchronisation strategy, and correlator configuration, were derived from first principles and validated experimentally. The system also demonstrates phase-coherent operation using a shared 10 MHz reference and a pulse-per-second

timing signal distributed across two USRP B210 units, providing a practical foundation for future scaling to larger arrays. The interferometer was validated through solar drift-scan observations at multiple east–west baselines. Interferometric fringes were detected successfully in all four observing sessions, with measured fringe periods agreeing with Earth-rotation predictions to within 16%.

1.4 Research Objectives

The objectives of this project are as follows:

- Design and build DIY pyramidal horn antennas for operation at 1.42 GHz, and study their gain, beam pattern, and impedance matching using analytical calculations and electromagnetic simulations.
- Develop and test front-end signal chains using low-noise amplifiers and filters, and identify a configuration that minimises the overall system noise figure using the Friis formula.
- Implement a phase-coherent two-element FX correlator using the USRP B210 and GNU Radio, and verify its performance through controlled bench experiments, including an extension to a four-element setup.
- Validate the interferometer through on-sky observations by detecting interferometric fringes and comparing their behaviour with predictions based on Earth rotation.
- Document the complete system, including design choices, calibration steps, and observation methods, in sufficient detail to support reproduction and further development.

1.5 Thesis Organisation

The remainder of this thesis is organised as follows. Chapter 2 outlines the theoretical foundations of radio interferometry, including the key principles underlying signal correlation and fringe formation. Chapter 3 presents the hardware design, covering the antenna, simulations, and front-end signal chain, with attention to noise performance. Chapter 4 describes the development of the digital backend, from the initial RTL-SDR setup to the USRP B210-based implementation and its extension to a four-element system for solar observations. Chapter 5 discusses the experimental results and their agreement with theoretical expectations, while Chapter 6 summarises the main findings, highlights the system's limitations, and suggests directions for future work.

Chapter 2

Foundations of Radio Interferometry

This chapter introduces the basic theory of radio interferometry. It begins with the response of a single antenna to incoming celestial radio waves and then extends to the two-element interferometer. Particular attention is given to what an interferometer measures, the physical meaning of that measurement, and how Earth's rotation improves observational capability.

2.1 Single-Dish Radio Astronomy

The most basic form of a radio telescope consists of a single antenna connected to a receiver. Examining the capabilities and limitations of such a system provides the necessary context for understanding why interferometry became essential in radio astronomy.

2.1.1 Collecting Radio Radiation

A radio telescope operates by intercepting electromagnetic radiation from the sky over an effective collecting area $A_{\text{eff}}$and converting the received energy into an electrical signal that can be measured by the receiver. The power collected from a source of flux density S (measured in Jansky, where 1 Jy = 10^{-26} W m^{-2} Hz^{-1}) over a bandwidth $\Delta\nu$ is given by:

$$P = \frac{1}{2} A_{\text{eff}} \cdot S \cdot \Delta\nu \tag{2.1}$$

The factor of $\frac{1}{2}$ arises because a single linearly polarised antenna responds to only one of the two orthogonal polarisations present in the incident wave [1]. The effective collecting area is related to the antenna gain G through:

$$A_{\text{eff}} = \frac{G\lambda^2}{4\pi} \tag{2.2}$$

For the horn antennas used in this project, whose gain was measured and simulated in Chapter 3 to be approximately 14.2 dBi at 1.42 GHz, this gives $A_{\text{eff}} \approx 0.12$ m^2. Although this collecting area is modest, it is sufficient for detecting a strong radio source such as the Sun.

2.1.2 Angular Resolution

An antenna does not receive radiation equally from all directions. Instead, it is characterised by a *primary beam*, which describes the directional sensitivity of the antenna. The response is strongest along the main lobe and decreases away from it. The angular width of the main lobe determines the finest angular detail that the telescope can distinguish. This is commonly described by the *half-power beamwidth* (HPBW), defined as the angle at which the sensitivity falls to half of its peak value. For an aperture of diameter D at wavelength λ [3]:

$$\theta_{\text{HPBW}} \approx \frac{\lambda}{D} \tag{2.3}$$

For the horn antenna used in this project ($D \approx 0.43$ m, $\lambda = 0.211$ m):

$$\theta_{\text{HPBW}} \approx \frac{0.211}{0.43} \approx 0.49 \text{ rad} \approx 28^\circ \tag{2.4}$$

This beam is extremely broad. Since the Sun subtends only about 0.5° on the sky, the horn cannot resolve its disc. Instead, it measures the total radio power received from the Sun as the source drifts through the beam. The received power at any instant is the integral of the sky brightness weighted by the beam pattern:

$$P(\nu) = \frac{A_{\text{eff}}}{2} \int I_\nu(\hat{s})\, A(\hat{s})\, d\Omega \tag{2.5}$$

where $I_\nu(\hat{s})$ is the sky brightness in direction $\hat{s}$ and $A(\hat{s})$ is the normalised beam pattern [6].A single-dish measurement therefore provides the total power within the beam, but does not reveal how that power is distributed within it or whether the source contains finer spatial structure.

2.1.3 The Radiometer Equation

Every electrical component in the receiver contributes thermal noise. This noise is added to the sky signal and ultimately sets the limit on how faint a source can be detected. The total noise in the system is conveniently represented by a single quantity, the *system temperature* T_{sys}, which is defined as the temperature of a matched resistor that would produce the same noise power:

$$T_{\text{sys}} = T_{\text{sky}} + T_{\text{ant}} + T_{\text{rec}} \tag{2.6}$$

Here, T_{sky} represents the contribution from the sky, T_{ant} accounts for antenna losses and ground pickup, and T_{rec} denotes the receiver noise, which is dominated in practice by the first low-noise amplifier (LNA) [3]. The minimum detectable flux density after integrating for a time t over a bandwidth $\Delta\nu$ is given by the *radiometer equation* [1]:

$$\boxed{S_{\text{min}} = \frac{2k_B T_{\text{sys}}}{A_{\text{eff}}\sqrt{\Delta\nu \cdot t}}} \tag{2.7}$$

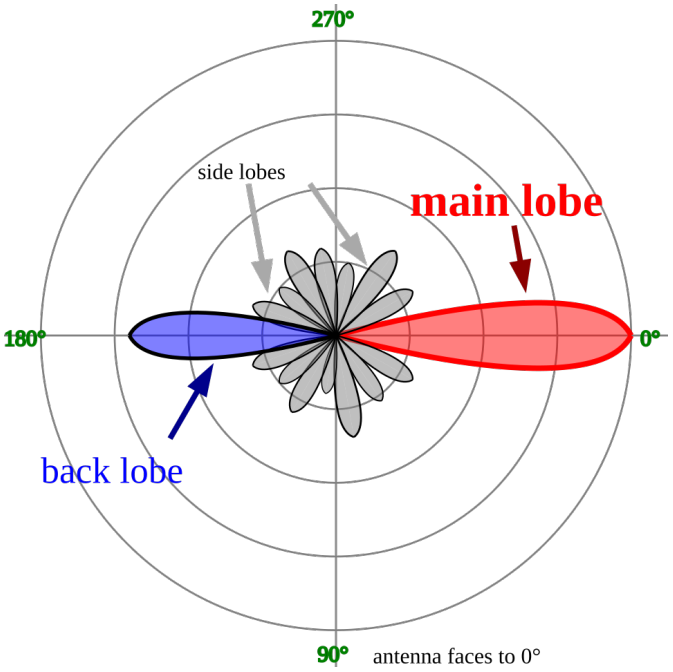

Figure 2.1: Radiation pattern of a directional antenna showing the main lobe, sidelobes, and back lobe. The angular width of the main lobe defines the effective beamwidth of the antenna, within which incoming radiation is integrated. The half-power beamwidth (HPBW) corresponds to the angular width of the main lobe at half maximum power. Image adapted from Wikimedia Commons [11].

The sensitivity improves with longer integration times and wider bandwidths. For example, quadrupling the integration time or doubling the bandwidth improves the sensitivity by a factor of $\sqrt{2}$. For our system ($T_{\text{sys}} \approx 60$ K, $A_{\text{eff}} \approx 0.12$ m^2, $\Delta\nu = 2$ MHz, $t = 30$ s):

$$S_{\text{min}} \approx \frac{2 \times 1.38 \times 10^{-23} \times 60}{0.12\sqrt{2 \times 10^6 \times 30}} \approx 2.4 \text{ Jy} \tag{2.8}$$

At 1.42 GHz, the Sun radiates at roughly 10^4–10^5 Jy [1], which is many thousands of times above our detection threshold. This gave us confidence from the beginning that the Sun would be a clear first target. We confirmed this expectation through single-dish tests before proceeding to interferometric observations.

2.1.4 Limitations of Single Dish

The main limitation is angular resolution. No matter how sensitive a single dish may be, it cannot resolve structures smaller than $\theta \approx \lambda/D$. To achieve a resolution of 1 arcsecond at a wavelength of 21 cm, which is comparable to that of a modest optical telescope, one would require a dish nearly 50 km in diameter. Even the 100-metre Green Bank Telescope, the world's largest fully steerable dish, can resolve only structures larger than about 7 arcminutes at this wavelength [1]. As described in Chapter 1, the solution is interferometry. By combining multiple smaller antennas and processing their signals together, the resolution is determined by the distance

between the antennas rather than by the size of the individual dishes.

2.2 Coordinate Systems

To determine the position of a source in the sky and predict when it will pass through our beam, two coordinate systems are required. The first is fixed relative to the stars, and the second is fixed relative to the ground.

2.2.1 Equatorial Coordinates: Right Ascension and Declination

The *equatorial coordinate system* is the standard framework used to specify positions on the sky. It is fixed to the celestial sphere and does not depend on the observer's location on Earth. Any position in this system is defined by two angles:

- **Right Ascension (α, RA):** The east-west coordinate, measured eastward from the vernal equinox in hours (0 h to 24 h). One hour corresponds to 15°.
- **Declination (δ, Dec):** The north-south coordinate, measured in degrees from the celestial equator (−90° to +90°).

For most celestial sources, RA and Dec remain fixed as the Earth rotates. The Sun, however, moves along the ecliptic over the course of the year. Its declination is approximately $\delta \approx -23°$ in December and $\delta \approx 0°$ at the March equinox [1]. This seasonal variation is important because it changes the fringe period that we expect to measure.

A closely related quantity is the *hour angle* H, which tells us how far a source has moved west of the local meridian:

$$H = \mathrm{LST} - \alpha \tag{2.9}$$

where LST is the Local Sidereal Time. At transit, when a source crosses the meridian and reaches its highest point in the sky, we have $H = 0$.

2.2.2 Horizontal Coordinates and Converting Between Systems

The *horizontal (Alt-Az) system* describes positions relative to the observer's horizon and is therefore useful for physically pointing the antennas. For an observer at latitude ϕ (positive north), the altitude of a source with declination δ and hour angle H is given by [6]:

$$\sin a = \sin\delta \sin\phi + \cos\delta \cos\phi \cos H \tag{2.10}$$

At IIT Indore ($\phi = 22.7°$N), the Sun at $\delta = 0°$ transits at an altitude of $a = 90° - 22.7° = 67.3°$. Our antennas are pointed south at the elevation that places the Sun within the beam at transit.

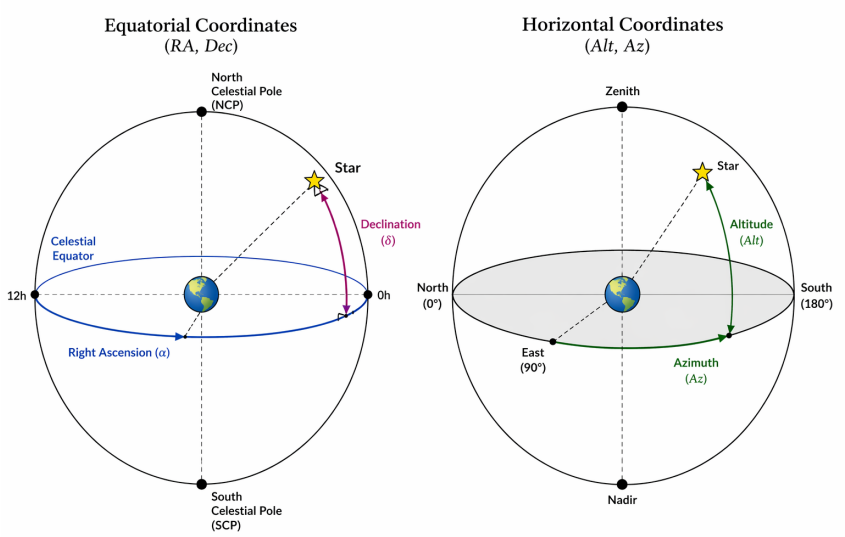

Figure 2.2: The equatorial (RA, Dec) and horizontal (Alt, Az) coordinate systems. As the Earth rotates, a source's RA and Dec remain fixed, while its altitude and azimuth change continuously

Because the antennas are fixed and the observations are carried out in drift scan mode, these equations allow us to predict exactly when the Sun enters and exits the beam. This defines the useful observing window for each session.

2.3 The Two-Element Interferometer

We now turn to the mathematical description of the interferometer. In simple terms, the signals received by two antennas are combined and compared to identify the part they have in common. This common component comes from the sky, while the remaining differences are mainly due to noise. Cross-correlation is the mathematical tool used to separate these contributions.

2.3.1 Geometric Delay

Consider two antennas located at positions $\mathbf{r}_1$ and $\mathbf{r}_2$, separated by the *baseline vector* $\mathbf{b} = \mathbf{r}_2 - \mathbf{r}_1$. A plane wave arriving from direction $\hat{s}$ reaches the two antennas at different times, because one antenna lies farther along the direction of propagation. This difference in arrival time is called the *geometric delay*:

$$\tau_g = \frac{\mathbf{b} \cdot \hat{s}}{c} \tag{2.11}$$

This quantity, which is simply a path difference divided by the speed of light, connects the physical separation of the antennas to the direction of the source in the sky [6]. It forms the basis for all of the interferometric relations that follow.

2.3.2 The Visibility Function

The signal received at each antenna from a cosmic source is not a pure sine wave, but broadband noise. Astronomical radio sources emit incoherently over a wide range of frequencies, and the received signal initially appears similar to thermal noise. The important point is that the *same* sky noise reaches both antennas, with a delay of τ_g, whereas the receiver noise in the two antennas is independent.

When the signals from the two antennas are multiplied and averaged over time, the independent noise terms average to zero, because positive and negative fluctuations occur randomly and without correlation. The common sky signal, however, remains and produces a nonzero result. This quantity is called the *complex visibility*:

$$\mathcal{V} = \left\langle V_1(t)\, V_2^*(t - \tau_g) \right\rangle \tag{2.12}$$

where the angle brackets denote time averaging [12]. The visibility is a complex number. Its amplitude $|\mathcal{V}|$ indicates the brightness of the source on the angular scale set by the baseline length, while its phase $\arg(\mathcal{V})$ gives information about the position of the source within the fringe pattern.

2.3.3 The van Cittert–Zernike Theorem: Connecting Fringes to Sky

So far, we have seen that the visibility contains information about the sky. The van Cittert–Zernike theorem gives the precise relationship [6, 12].

The sky can be thought of as being made up of many small patches, each emitting independently. For most astronomical sources, this is an excellent approximation. Each patch contributes to the visibility with a phase determined by its position relative to the baseline direction. When all of these contributions are added together, the result is an integral over the sky brightness distribution. If sky positions are expressed using direction cosines (l, m), and the baseline is written in units of wavelength as $(u, v) = \mathbf{b}/\lambda$, the theorem reduces, for a compact field of view, to a two-dimensional Fourier transform [12]:

$$\boxed{\mathcal{V}(u, v) = \iint I(l, m)\, e^{-i2\pi(ul+vm)}\, dl\, dm} \tag{2.13}$$

This is one of the central equations of aperture synthesis. It shows that *the visibility measured by the interferometer is the Fourier transform of the sky brightness distribution*. Each baseline samples one spatial frequency component of the sky, in the same way that a single Fourier coefficient represents one frequency component of a signal. If a sufficient number of baselines are measured, the sky image can, in principle, be reconstructed through an inverse Fourier transform.

2.4 The UV Plane and Aperture Synthesis

2.4.1 The UV Plane

The baseline coordinates (u, v), measured in units of the observing wavelength, define a point in a two-dimensional Fourier space known as the *uv plane*. Each visibility measurement adds one point, along with its complex conjugate, to this plane. Building up coverage of the uv plane through multiple baselines or through observations over time is what makes image reconstruction possible.

For an east-west baseline of physical length B, the uv coordinates depend on the source declination δ, hour angle H, and the observer latitude ϕ [6]:

$$u = \frac{B}{\lambda} \cos\delta \sin H \tag{2.14}$$

$$v = \frac{B}{\lambda} (\sin\delta \cos\phi - \cos\delta \cos H \sin\phi) \tag{2.15}$$

2.4.2 Aperture Synthesis and Earth-Rotation

Aperture synthesis becomes particularly effective when the rotation of the Earth is taken into account. At any instant, a single baseline samples only one point in the uv plane. As the Earth rotates, the source hour angle H changes, so the orientation of the baseline with respect to the sky also changes. As a result, the same baseline traces an *elliptical arc* in the uv plane. This effect is known as *Earth-rotation aperture synthesis*. Over the duration of an observation, a single physical baseline can therefore sample many points in the uv plane, allowing the array to behave as if it were more densely populated than its physical layout would suggest [6]. As the Earth rotates, a single east-west baseline traces an elliptical arc in the uv plane.

2.5 The Fringe Pattern

The geometric delay between the two antennas changes continuously as the Earth rotates. For a signal at frequency ν, this changing delay causes the visibility phase to vary according to $\phi = 2\pi\nu\tau_g$. As a result, the correlator output shows a slow sinusoidal variation as the source moves across the sky, producing what is known as the *fringe pattern*. Detecting these fringes is the main observational goal of this project.

For an east-west baseline B observing a source at declination δ and hour angle $H(t)$ [6]:

$$\mathcal{V}(t) = A\, e^{\,i\phi(t)}, \quad \phi(t) = 2\pi \frac{B}{\lambda} \cos\delta \sin H(t) \tag{2.16}$$

2.5.1 Fringe Period

Near transit ($H \approx 0$), the phase varies almost linearly with time. The time taken for one complete oscillation is called the *fringe period*:

$$\boxed{T_{\text{fringe}} = \frac{\lambda}{B \cos\delta \cos H \cdot \omega_\oplus}} \tag{2.17}$$

where $\omega_\oplus = 7.2921 \times 10^{-5}$ rad s^{-1} is Earth's rotation rate. In practical units, noting that $\omega_\oplus = 7.2921 \times 10^{-5}$ rad s^{-1} $= 0.25^\circ$ min^{-1}, this may be written as:

$$T_{\text{fringe}} = \frac{(\lambda/B) \cdot (180/\pi)}{0.25 \cos\delta \cos H} \quad \text{minutes} \tag{2.18}$$

The fringe period is shorter for longer baselines, while sources at higher declination, that is, farther from the celestial equator, have longer fringe periods. The quantity T_{fringe} also depends on $\cos H$, so it is not strictly constant during a drift scan. It is smallest near transit and increases as the source moves away from the meridian. Therefore, the values measured during observation should be interpreted as averages over the corresponding range of hour angle.

2.5.2 Fringe Amplitude and Source Structure

The visibility amplitude $|\mathcal{V}|$ is not just a fitting parameter; it also carries direct information about the angular size of the source. If the source is sufficiently extended, radiation from different parts of it reaches the two antennas with different phases. These contributions partly cancel each other, which reduces the fringe amplitude. In contrast, if the source is smaller than the fringe spacing, approximately λ/B in angular terms, radiation from across the source arrives nearly in phase, and the measured amplitude simply corresponds to the total flux.

For a uniform circular disk of angular diameter θ_s, the visibility amplitude is given by [6]:

$$|\mathcal{V}(u)| = S \cdot \left| \frac{2J_1(\pi\theta_s u)}{\pi\theta_s u} \right| \tag{2.19}$$

where J_1 is the first-order Bessel function. The first zero in the amplitude occurs at a baseline length of $B_{\text{null}} = 1.22\lambda/\theta_s$. For the Sun ($\theta_s = 0.5^\circ = 8.73 \times 10^{-3}$ rad):

$$B_{\text{null}} = \frac{1.22 \times 0.211}{8.73 \times 10^{-3}} \approx 29.5 \text{ m} \tag{2.20}$$

All of our baselines (1.00–2.40 m) are much shorter than this value. The Sun is therefore *completely unresolved* on every baseline used in this study, and the fringe amplitude is expected to remain essentially constant across all baselines. This provides a clear and testable prediction, which is examined directly in Chapter5.

2.6 Mathematical Tools: Fourier Transform, Correlation, and Convolution

Three mathematical operations lie at the core of radio interferometry, both in theory and in the signal-processing pipeline described in Chapter 4. These are the Fourier transform, cross-correlation, and convolution. Although they are often introduced in a highly mathematical way, their basic roles are quite intuitive. The Fourier transform tells us which frequencies are present in a signal. Cross-correlation tells us how similar two signals are and whether one is delayed relative to the other. Convolution describes how a signal is modified when it passes through a system that spreads or blurs it. Together, these operations allow a faint astronomical signal to be identified within noisy measurements from two antennas and, ultimately, make it possible to reconstruct information about the sky.

2.6.1 The Fourier Transform

The Fourier transform expresses a signal in terms of its frequency components, representing it as a sum of sinusoids with different frequencies.

Formally, a time-domain signal $v(t)$ can be written in the frequency domain as [1]:

$$\tilde{V}(\nu) = \int_{-\infty}^{\infty} v(t)\, e^{-i2\pi\nu t}\, dt, \qquad v(t) = \int_{-\infty}^{\infty} \tilde{V}(\nu)\, e^{+i2\pi\nu t}\, d\nu \tag{2.21}$$

The first expression converts a time-domain signal into its frequency-domain form, $\tilde{V}(\nu)$, while the second reconstructs the original signal from that frequency representation. These are not different signals, but two equivalent descriptions of the same information.

In practice, the signals used in this project are not continuous. The analogue-to-digital converter in the backend samples the incoming voltage at discrete time intervals. For N discrete samples, the Discrete Fourier Transform (DFT) is given by

$$\tilde{V}[k] = \sum_{n=0}^{N-1} v[n]\, e^{-i2\pi nk/N} \tag{2.22}$$

This produces $N/2$ positive-frequency channels, each with width $\Delta\nu = f_s/N$, where f_s is the sampling rate. For our USRP B210 backend ($f_s = 2.0$ MSPS, $N = 8192$), each channel has a frequency resolution of approximately 244 Hz. This resolution is high enough to isolate narrow radio-frequency interference (RFI) features and reduce their effect without significantly disturbing the hydrogen-line signal.

Computing the DFT directly from its definition requires $O(N^2)$ operations, which becomes impractical when N is large. The Fast Fourier Transform (FFT) reduces this to $O(N \log N)$ by reorganising the computation [12]. This makes real-time spectral processing practical even on a standard laptop.

Windowing: Keeping RFI in its Lane

When a finite block of N samples is transformed using the FFT, the calculation implicitly assumes that the signal repeats periodically beyond the boundaries of that block. Real signals do not satisfy this condition exactly. As a result, a mathematical artefact known as *spectral leakage* appears. In this effect, power from a strong signal in one frequency channel spreads into adjacent channels and can obscure a weaker nearby signal.

To reduce this problem, the data are *windowed* before the FFT is applied. This means that each block of samples is multiplied by a smooth tapering function that brings the signal closer to zero at both ends. In our correlator, a Blackman–Harris window is used. This suppresses sidelobes to below -92 dB, allowing signals that are far weaker than nearby RFI sources to remain measurable. The ZCU216 RFSoC backend improves further on this by using a Polyphase Filter Bank (PFB), which achieves sidelobe suppression better than -100dB, as described in Chapter4.

Role of the Fourier Transform in Interferometry

The importance of the Fourier transform in interferometry is deeper than its computational usefulness. The van Cittert–Zernike theorem (Eq.2.13), derived in Section2.3.3, shows that the visibility measured by an interferometer is the Fourier transform of the sky brightness distribution. Recovering an image of the sky from measured visibilities is therefore, in principle, an inverse Fourier transform. The FFT makes this inversion computationally feasible. This relationship lies at the heart of radio interferometry and is one of the key reasons for its power as an observational technique [1].

2.6.2 Cross-Correlation

Cross-correlation is the mathematical operation that underlies interferometry. In simple terms, it helps identify the part of two noisy signals that is common to both. Since both antennas are pointed at the same region of the sky, they receive the same astronomical signal, but with a small difference in arrival time because the wave reaches one antenna slightly before the other. At the same time, each antenna also receives contributions from thermal receiver noise, local radio-frequency interference, and other random fluctuations.These unwanted contributions are usually independent in the two signal chains. Cross-correlation combines the two signals and averages the result over time. Since the noise is uncorrelated, its average contribution tends to zero. The sky signal, however, is coherent and appears in both antennas with the same underlying structure, so it survives the averaging process and produces a nonzero result. This rejection of uncorrelated noise is the fundamental reason an interferometer can detect signals that are much weaker than the noise level in either individual antenna.

Mathematically, the cross-correlation of signals $V_1(t)$ and $V_2(t)$ as a function of time lag τ is given by [12]:

$$R_{12}(\tau) = \lim_{T\to\infty} \frac{1}{2T} \int_{-T}^{T} V_1(t)\, V_2^*(t-\tau)\, dt \tag{2.23}$$

For two independent noise signals, $\langle n_1(t), n_2^*(t-\tau)\rangle = 0$, so uncorrelated noise averages to zero. The coherent sky signal, by contrast, remains after averaging and appears as a nonzero cross-correlation at the geometric delay $\tau = \tau_g$. The amplitude of this peak contains information about the source brightness, while its phase indicates the source position within the fringe pattern.

The Cross-Correlation Theorem and the FX Architecture

An important result links cross-correlation in the time domain to multiplication in the frequency domain [12]:

$$R_{12}(\tau) \;\longleftrightarrow\; \tilde{V}_1(\nu) \cdot \tilde{V}_2^*(\nu) \tag{2.24}$$

This theorem forms the basis of the *FX correlator architecture* used in both of our backends. Rather than computing the cross-correlation integral directly in the time domain, which would require a separate calculation for every possible delay τ, we first Fourier transform both signals, which is the **F** stage, and then multiply them channel by channel in the frequency domain, which is the **X** stage. The result is mathematically equivalent to the full cross-correlation, but far more efficient to compute. On a standard laptop running GNU Radio with our USRP B210, this method processes 2.4 million samples per second across 8192 frequency channels in real time, which would be impractical with direct time-domain correlation.

2.6.3 Convolution

Convolution describes how a signal or image is modified when it passes through a system that spreads or blurs it. A familiar example is a photograph taken slightly out of focus, where the recorded image is the true scene convolved with the blurring pattern of the optical system. In radio interferometry, the same idea applies. A map formed from incomplete interferometric data is the convolution of the true sky brightness with the *dirty beam*, which represents the response imposed by the limited and incomplete set of measured baselines.

Formally, the dirty image is the convolution of the true sky brightness $I(l, m)$ with the dirty beam $B(l, m)$ [12]:

$$I^D(l,m) = \iint I(l',m')\, B(l-l', m-m')\, dl'\, dm' \equiv (I * B)(l,m) \tag{2.25}$$

The *convolution theorem* explains why this blurring occurs. It states that convolution in the image domain corresponds to multiplication in the Fourier domain:

$$I^D = I * B \longleftrightarrow \mathcal{V}(u, v) \cdot S(u, v) \tag{2.26}$$

where $S(u, v)$ is the sampling function, equal to one at every measured (u, v) point and zero elsewhere. This relation shows clearly why incomplete uv coverage leads to a dirty image. In the Fourier domain, the true visibility is multiplied by the sparse sampling function $S(u, v)$. In the image domain, this becomes a convolution of the true sky brightness with the corresponding dirty beam [12]. Deconvolution methods such as CLEAN are designed to reverse this effect and recover an estimate of the true brightness distribution $I(l, m)$ from $I^D(l, m)$.

For our interferometer, full imaging is beyond the scope of the present work. Our four baselines sample only a narrow arc in the uv plane, which is insufficient for reconstructing a reliable sky map. Convolution still plays an important role. It explains why the measured fringe amplitude already reflects the Sun's angular size (Section 2.5.2) and also provides the mathematical basis for any future image reconstruction with this instrument.

Chapter 3

Hardware Design: Antenna and Front-End Electronics

This chapter describes the hardware used in the system, including the antennas that receive radio signals from the sky and the electronics that prepare those signals for digitisation. It also outlines the reasoning behind the main design choices made in this work.

The design philosophy throughout this project is transparency. Every component was selected not only for its performance, but also so that a student can understand its function, measure its behaviour independently, and replace or modify it if necessary. This is one of the main features that distinguishes an educational instrument from a professional one. Each antenna chain follows the same path: a pyramidal horn antenna converts incoming electromagnetic radiation into an electrical signal within a waveguide; a low-noise amplifier increases that signal before additional noise can accumulate; a pair of filters removes interference; and a second amplifier stage raises the signal to a level that the digital backend can digitise. Two identical chains feed one digital backend (Chapter 4), which cross-correlates their outputs to produce interferometric visibility data.

3.1 System Overview

Figure 3.1 shows the overall architecture. shows the overall architecture. Each of the n antenna elements feeds its own identical front-end module, which in turn feeds its own synchronised SDR or RFSoC backend. All backends send their digitised, channelised data to a central laptop, where cross-correlation is carried out in software and the interferometric fringes are extracted.

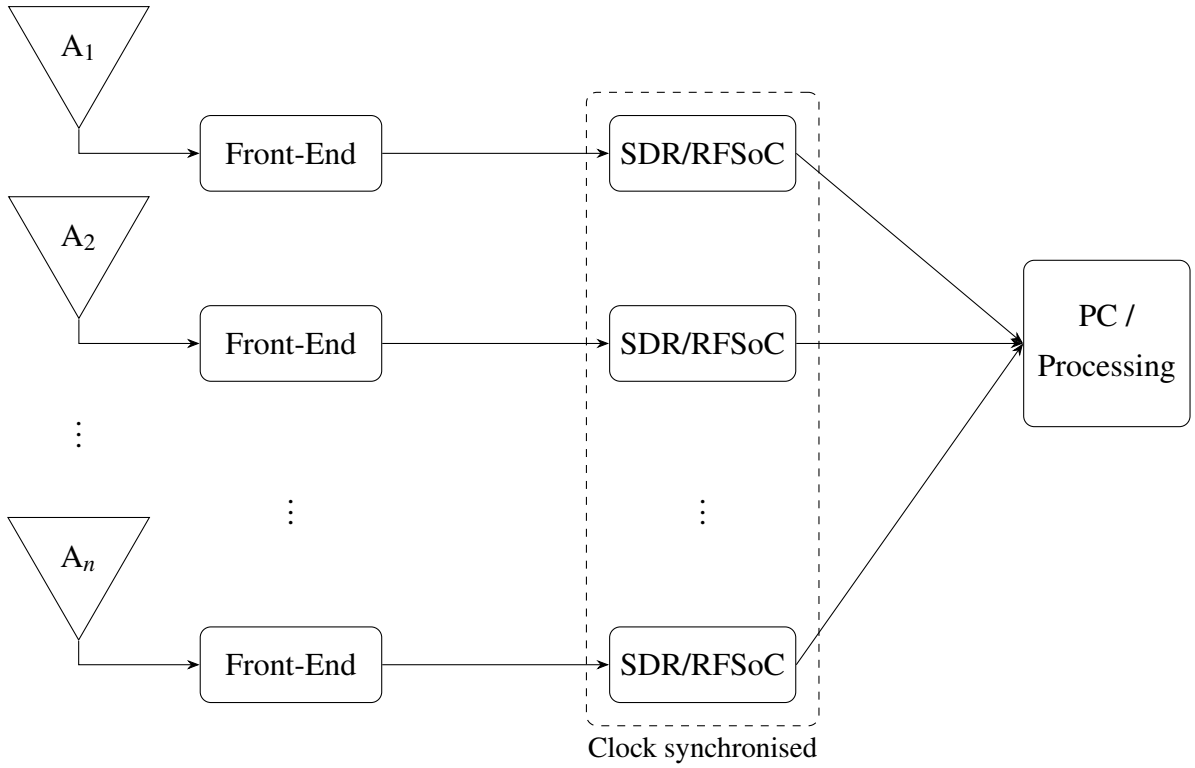

Figure 3.1: Block diagram of the n-element instructional radio interferometer. Each antenna feeds an identical front-end module and a synchronised SDR/RFSoC backend. All backends connect to a central laptop for software correlation and fringe analysis. The dashed box highlights the clock-synchronisation requirement that motivated the choice of digital backend, as discussed in Chapter 4.

The modular structure offers a clear educational advantage: each block can be characterised independently before the complete chain is assembled. The LNA gain can be measured with a vector network analyser, the filter passband can be verified, and the ADC clock coherence can be tested before any astronomical observation is attempted.

3.2 Antenna Design and Characterisation

The antenna is the first element in the signal chain and plays a central role in determining what enters the system. Any signal not received at this stage cannot be recovered later through amplification. The system is designed to operate at $f_0 = 1.420$ GHz ($\lambda_0 = 0.211$ m), corresponding to the neutral hydrogen line. This is the frequency at which neutral hydrogen emits radiation through a change in the electron spin state. It is one of the most important spectral lines in radio astronomy and is protected under international regulations, which helps keep this part of the spectrum relatively free from interference.

For this instrument, *pyramidal horn antennas* were chosen. A horn antenna is essentially a waveguide whose walls flare outward into a pyramidal shape. This geometry improves the transition between the waveguide and free space, making radiation more efficient than from a

simple waveguide opening. It also defines the angular region of the sky from which radiation is received, giving the antenna a predictable gain and beam pattern [13].

From a practical point of view, pyramidal horns are well suited to this project. They can be built from sheet metal using relatively simple tools, which makes them a good choice for a student-led instrument. Their design is supported by well-established analytical formulas, and their performance can be verified independently through electromagnetic simulation.

3.2.1 Design Goals

Before fabrication, a few key requirements were set for the antenna design:

- The antenna should operate at a centre frequency of $f_0 = 1.420\,\text{GHz}$, with good impedance matching near this frequency.
- A gain of around 10 dBi was targeted, which is sufficient for detecting solar emission while keeping the antenna small enough for rooftop use.
- The waveguide section should support only the dominant TE_{10} mode at 1.42 GHz. Higher-order modes would produce unwanted radiation patterns and make the antenna behaviour less predictable.
- The design should be compatible with the sheet-metal fabrication facilities available at IIT Indore.

3.2.2 Analytical Design Equations

The key equations for a pyramidal horn are taken from standard antenna theory [13]. These equations were used both to design the compact horn and to assess the laboratory horn that was ultimately used.

Gain from aperture area. The gain of a pyramidal horn relative to an isotropic radiator is:

$$G = \frac{4\pi\, A_{\text{ap}}}{\lambda^2}\,\eta \tag{3.1}$$

where $A_{\text{ap}} = W_a \times H_a$ is the physical aperture area (W_a = aperture width, H_a = aperture height), λ is the wavelength, and η is the *aperture efficiency*, which is the fraction of the physical aperture area that effectively contributes to radiation. For a pyramidal horn, η typically lies in the range 0.5–0.7; the optimum design value is $\eta \approx 0.511$ [13]. In this work, we adopt $\eta = 0.65$, which is appropriate for the laboratory horn because its flare length provides a good phase front across the aperture. The CST simulation results in Section 3.2.5 support this choice.

Waveguide cutoff frequency. Inside the rectangular throat, radio waves propagate in the dominant TE_{10} mode only above the cutoff frequency:

$$f_c = \frac{c}{2\,a_{\text{th}}} \tag{3.2}$$

where a_{th} is the broader inner dimension of the rectangular throat [13]. For the antenna to operate at 1.42 GHz, we require $f_c < 1.42$ GHz. At the same time, we require $f_c > 0.71$ GHz to prevent the next higher mode, TE20, from propagating at the operating frequency. In practice, choosing ath such that f_c lies between 0.7 and 0.9 times the operating frequency provides good single-mode performance.

Half-power beamwidths. The angular width of the main beam in each plane is determined by the corresponding aperture dimension:

$$\text{HPBW}_W \approx \frac{\lambda}{W_a}\ (\text{rad}), \qquad \text{HPBW}_H \approx \frac{\lambda}{H_a}\ (\text{rad}) \tag{3.3}$$

These expressions are approximate and are valid when the aperture is several wavelengths across [13]. A larger aperture produces a narrower beam and higher gain, whereas a smaller aperture produces a broader beam. In drift-scan observations, where the antenna remains fixed and the source moves across the field of view, a slightly wider beam can be advantageous.

3.2.3 10 dBi Compact Design

Starting from the gain target of $G = 10$ dBi ($G_{\text{lin}} = 10$), the required aperture can be estimated as follows:

$$A_{\text{required}} = \frac{G_{\text{lin}}\,\lambda^2}{4\pi\eta} = \frac{10 \times (0.211)^2}{4\pi \times 0.65} \approx 0.0546\ \text{m}^2 \tag{3.4}$$

Choosing an aperture aspect ratio of $W_a : H_a \approx 1.4$, which is close to the standard rectangular waveguide throat ratio of 2 : 1:

$$W_a \approx 0.28\ \text{m}, \qquad H_a \approx 0.195\ \text{m} \tag{3.5}$$

The minimum throat dimension required for TE_{10} propagation at 1.42 GHz is:

$$a_{\text{th,min}} = \frac{c}{2f_0} = \frac{3 \times 10^8}{2 \times 1.42 \times 10^9} \approx 0.106\ \text{m} \tag{3.6}$$

A practical choice is $a_{\text{th}} = 0.12$ m, $b_{\text{th}} = 0.06$ m (maintaining the 2:1 ratio), with a flare length of $L \approx 0.3$–0.4 m. This compact horn was developed as the target design specification. In practice, however, the larger laboratory horns described below were used, as a matched pair was already available. A matched pair of laboratory horns was already available in the department,

offering higher gain and a more comfortable single-mode margin, so fabrication of the compact design was deferred.

3.2.4 Horn Antenna: Dimensions and Predicted Performance

For all outdoor observations presented in this thesis, a pair of matched laboratory horns already available in the department was used. These horns exceed the 10 dBi design target, which improves sensitivity. Using an identical matched pair was also important for maintaining similar signal chains on both antennas.

The physical dimensions of each horn are:

$$W_a = 0.435 \text{ m}, \quad H_a = 0.330 \text{ m}, \quad L = 0.920 \text{ m} \tag{3.7}$$

$$a_{\text{th}} = 0.155 \text{ m}, \quad b_{\text{th}} = 0.106 \text{ m} \tag{3.8}$$

Aperture area and predicted gain.

$$A_{\text{actual}} = 0.435 \times 0.330 = 0.1436 \text{ m}^2 \tag{3.9}$$

$$G_{\text{actual}} = \frac{4\pi \times 0.1436}{(0.211)^2} \times 0.65 \approx 26.2 \text{ (linear)} \approx 14.2 \text{ dBi} \tag{3.10}$$

Waveguide cutoff.

$$f_c = \frac{3 \times 10^8}{2 \times 0.155} \approx 0.97 \text{ GHz} \tag{3.11}$$

This value lies well below 1.42GHz, confirming single-mode TE_{10} operation, and well above 0.71GHz, which is half of 1.42 GHz, confirming that the TE_{20} mode is suppressed.

Half-power beamwidths.

$$\text{HPBW}_W = \frac{0.211}{0.435} \times \frac{180^\circ}{\pi} \approx 27.8^\circ, \qquad \text{HPBW}_H = \frac{0.211}{0.330} \times \frac{180^\circ}{\pi} \approx 36.7^\circ \tag{3.12}$$

The beam is relatively wide in both planes. For drift-scan observations, this is an advantage rather than a limitation. With Earth rotating at $0.25^\circ, \text{min}^{-1}$, the Sun remains within the 27.8° H-plane beam for approximately $27.8/0.25 \approx 111$ minutes around transit. This time interval is sufficient to capture multiple fringe cycles, even for the longest baseline used in this work.

Table 3.1compares the compact design target with the laboratory horn used in the final system.

Table 3.1: Comparison of the compact 10 dBi design target with the horn antenna used in all observations. Both are designed to operate at $f_0 = 1.420$ GHz.

Parameter	Compact Design	Our Horn Atenna
Aperture $W_a \times H_a$ (m)	0.28×0.195	0.435×0.330
Aperture area (m^2)	0.0546	0.1436
Predicted gain (dBi)	~10	~14.2
Throat a_{th} (m)	0.120	0.155
Cutoff f_c (GHz)	1.25	0.97
Flare length L (m)	0.30–0.40	0.920
HPBW (H/E plane, deg)	~35 / ~45	27.8 / 36.7

The compact design has a cutoff frequency of 1.25 GHz, which is relatively close to the operating frequency of 1.42 GHz. As a result, the antenna would still function, but it would operate near the edge of the single-mode region and would require careful waveguide throat fabrication. In comparison, the laboratory horn, with $f_c = 0.97$ GHz, offers a more comfortable margin and is therefore more tolerant of small fabrication errors.

3.2.5 Antenna Simulation in CST Studio Suite

Before relying on the analytical estimates, the laboratory horn was modelled in CST Studio Suite, a commercial electromagnetic simulation package that numerically solves Maxwell's equations for the full three-dimensional antenna geometry. The main aims were to verify the accuracy of the analytical predictions for the chosen horn dimensions and to generate the beam pattern.

The model includes all the main physical dimensions of the laboratory horn, including the waveguide throat, flare section, open aperture, and the coaxial-to-waveguide feed transition.

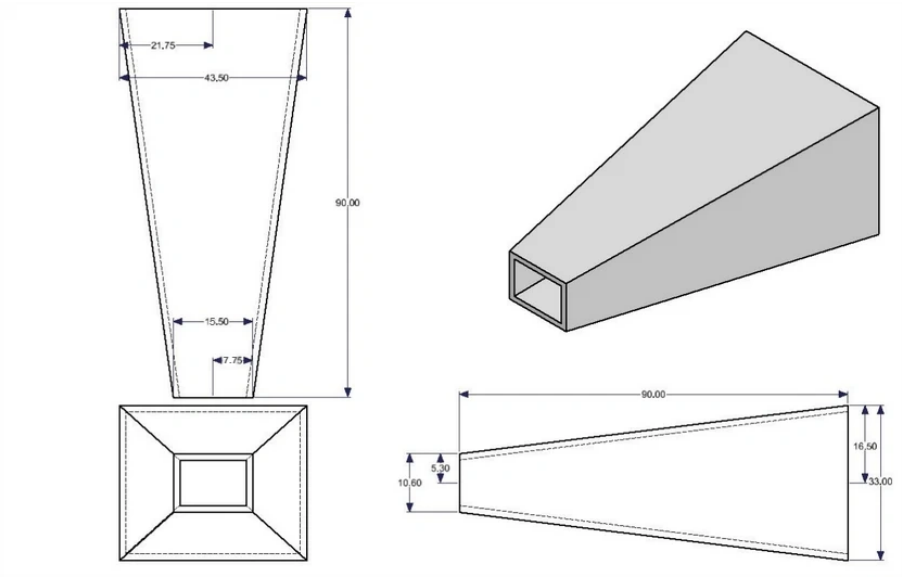

Figure 3.2: Flare section of the horn, showing the aperture width W_a, aperture height H_a, and flare length L. Together with the throat dimensions, these parameters fully define the antenna geometry.

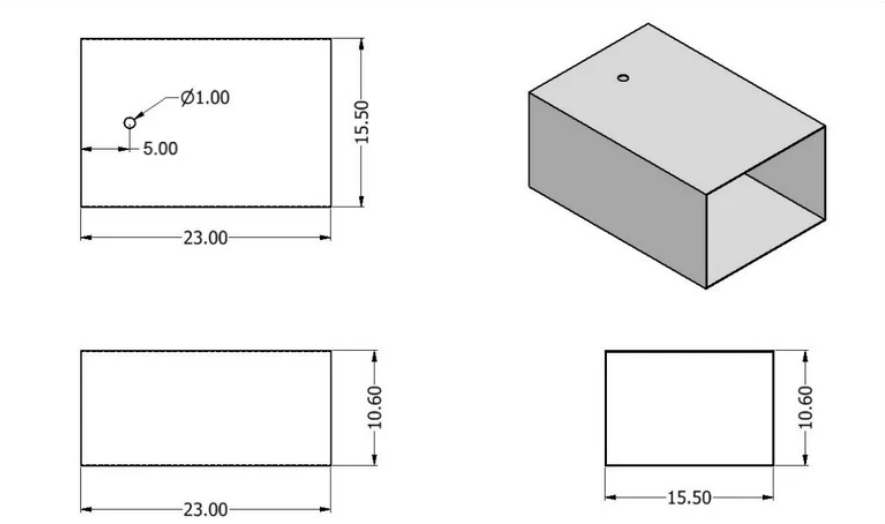

Figure 3.3: Waveguide throat and feed transition, showing the broad dimension a_{th}, narrow dimension b_{th}, and the feed point where the coaxial connector excites the TE_{10} mode.

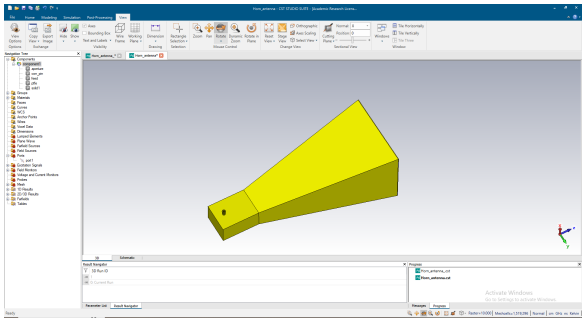

(a) Isometric view of the full horn model.

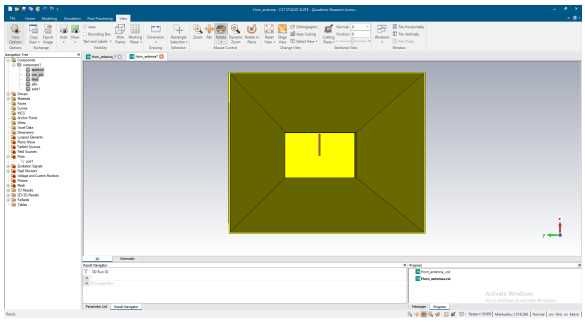

(b) Front view showing aperture dimensions.

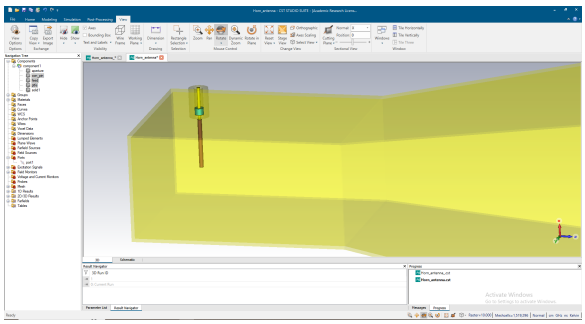

(c) Zoomed in View of feed.

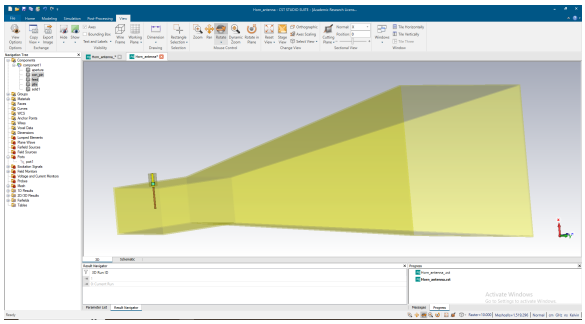

(d) Side view showing flare length.

Figure 3.4: Simulation model of the laboratory horn, illustrating multiple views of the geometry, including the aperture, flare section, and feed region.

Figure 3.5 shows the simulated return loss (S_{11}) as a function of frequency. The parameter S_{11} gives the fraction of input power reflected back toward the source. A value below -10 dB means that less than 10% of the power is reflected. At 1.42 GHz, the simulated minimum is approximately -9.22 dB, corresponding to a reflected power of about 12%. Although this is slightly above the conventional -10 dB guideline, the matching is still adequate for the present solar interferometry application, where the Sun provides a comparatively strong radio signal.



Figure 3.5: Simulated return loss S_{11} of the laboratory horn antenna. The minimum value of -9.22 dB at 1.42 GHz indicates good impedance matching near the design frequency.

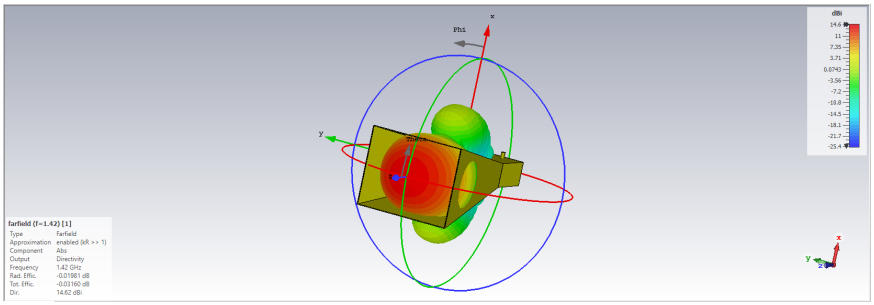

Figure 3.6: Simulated 3D gain pattern at 1.42 GHz. The peak gain of $\approx$ 14.2dBi agrees with the analytical prediction of Equation 3.10, validating the use of $\eta = 0.65$.

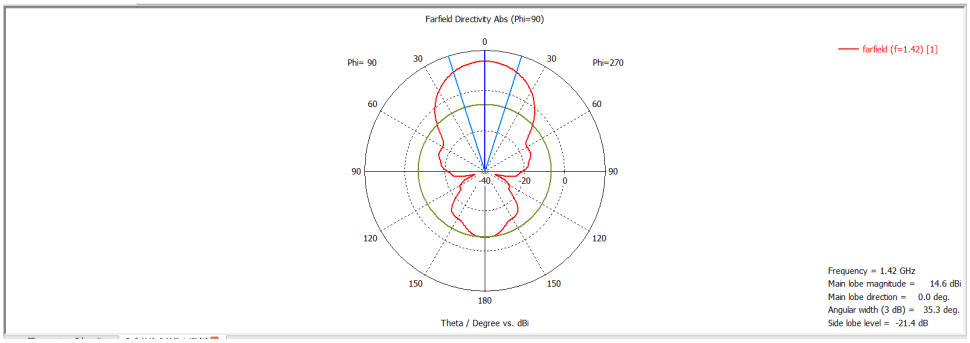

Figure 3.7: Simulated far-field radiation pattern (|Directivity|) at 1.42 GHz for the $\phi = 90°$ plane. The main lobe is centered at 0° with a peak directivity of 14.6dBi. The half-power beamwidth (3dB) is approximately 35.3°, and the side lobe level is −21.4 dB, consistent with the expected directional characteristics of the antenna.

The close agreement between the simulated gain pattern, the return loss dip, and the analytical predictions gives confidence that the horn performs as expected before being used for astronomical observations.

3.3 Front-End Electronics

After the antenna, the received signal is extremely weak. At 1.42 GHz, a typical sky temperature of $\sim$ 50–100 K corresponds to a noise power of order $k_B T \Delta\nu \sim 10^{-15}$–$10^{-14}$ W over a 2 MHz bandwidth. Signals at this level must be amplified before digitisation, while keeping added noise low and suppressing local interference.

The front-end design follows three main considerations. First, the noise performance is set mainly by the initial stage (Section 3.3.2), so a low-noise amplifier should be placed immediately after the antenna. Second, filtering is applied after the initial gain to avoid sensitivity loss from insertion loss. Third, the signal chain is kept simple, with each component serving a specific purpose.

3.3.1 Components

The front-end is built from four off-the-shelf RF modules selected for hydrogen-line observations at 1.42 GHz.

Low-Noise Amplifier: Nooelec SAWbird+ H1

The low-noise amplifier (LNA) is the most important component in the receiver chain because it provides high gain while adding very little noise. From the Friis relation (Eq. 3.15), the noise contribution of later stages is suppressed by the gain of the first stage, making a low-noise, high-gain LNA essential.

The Nooelec SAWbird+ H1 is designed for hydrogen-line observations at 1.420 GHz and includes a SAW (Surface Acoustic Wave) filter ahead of the amplifier. This filter suppresses strong out-of-band interference, such as mobile, Wi-Fi, and radar signals, helping to prevent saturation and maintain linear operation.



Figure 3.8: The Nooelec SAWbird+ H1 LNA. The module provides a gain of approximately 39.5 dB while contributing very little additional noise (NF $\approx$ 0.04 dB). Its integrated SAW pre-filter also helps suppress strong out-of-band interference before amplification.

At 1.42 GHz, the module provides a gain of approximately 39.5 dB with a noise figure of about 0.04 dB. It can be powered either through USB (5 V) or by a bias-tee that supplies DC power through the coaxial cable.

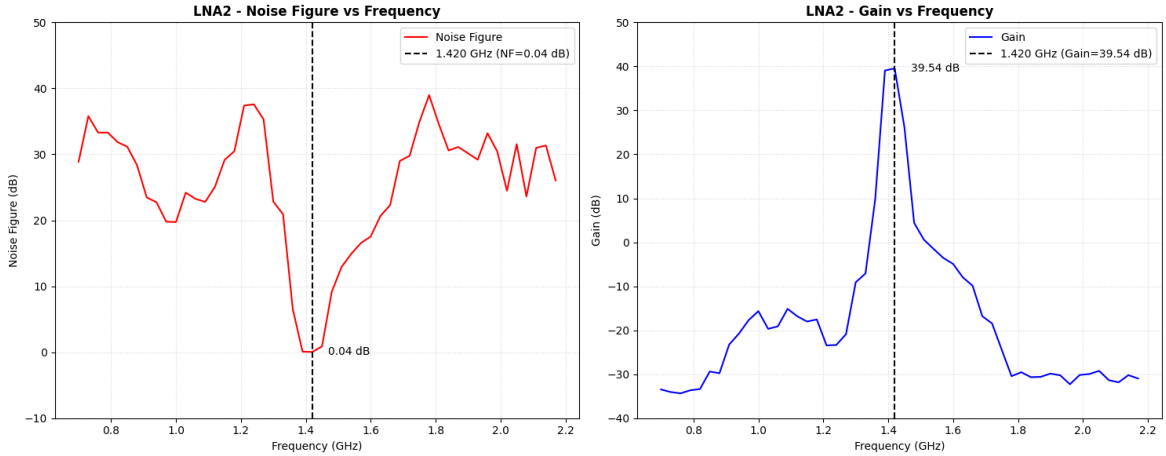

Figure 3.9: Measured gain and noise figure of the SAWbird+ H1 as functions of frequency. Both peak near 1.42 GHz, confirming that the module is well suited to hydrogen-line observations.

Band-Pass Filter: Mini-Circuits VBFZ-1400-S+

The band-pass filter (BPF) passes signals in the 1350–1450 MHz range and provides more than 40 dB of rejection outside this band. Its main role is to suppress residual out-of-band interference not removed by the SAW pre-filter, thereby protecting the second LNA stage and the ADC, while also limiting the noise bandwidth passed to the digital backend. At IIT Indore, major RFI sources include LTE cellular signals (700–2100MHz), Wi-Fi (2.4 and 5GHz), and radar emissions. After the ~39.5 dB gain of the first LNA, any remaining interference is also amplified, so the BPF attenuates this unwanted power before the second gain stage.



Figure 3.10: Mini-Circuits VBFZ-1400-S+ band-pass filter, which isolates the 1350–1450 MHz hydrogen-line window.

The filter has a passband of 1350–1450MHz, an in-band insertion loss of approximately 2dB, and an out-of-band rejection greater than 40 dB.

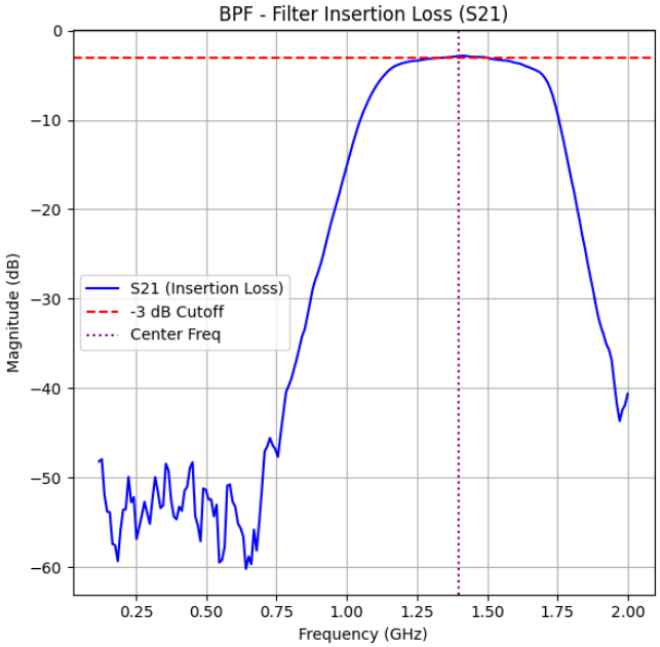

Figure 3.11: S-parameter response of the VBFZ-1400-S+. The 2dB in-band insertion loss and >40dB out-of-band rejection are clearly visible. Because this filter is placed after the first LNA, its 2dB loss contributes negligibly to the overall system noise figure (Section 3.3.2).

High-Pass Filter: Mini-Circuits SHP-1000+

The high-pass filter provides an additional layer of low-frequency rejection. FM radio broadcasts (88–108MHz), laboratory switching power supplies, and general electrical noise below 1GHz can, if they are strong enough, alias into the signal band during digitisation. The SHP-1000+ suppresses signals below roughly 1GHz with more than 40dB of rejection.

The filter has a cutoff frequency of approximately 1GHz, an in-band insertion loss of less than 1dB over 1–2GHz, and more than 40dB rejection below 500 MHz.



Figure 3.12: Mini-Circuits SHP-1000+ high-pass filter, placed at the end of the chain to block FM radio and other sub-GHz interference.

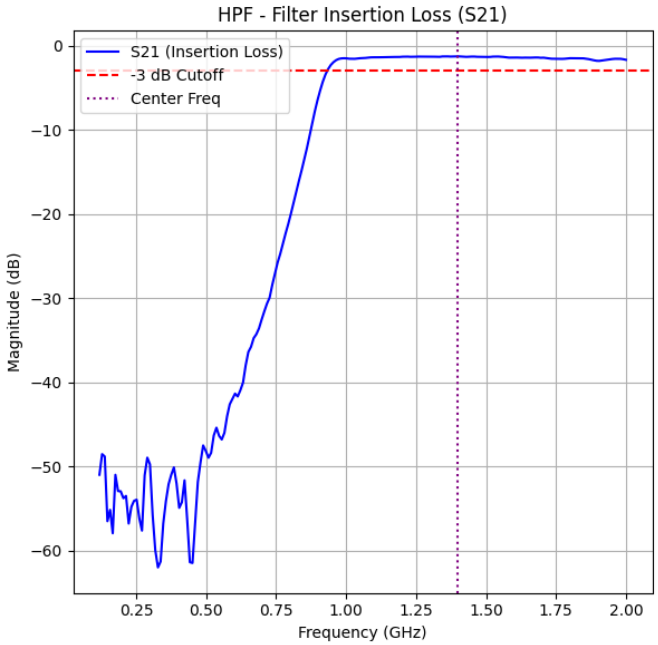

Figure 3.13: Measured response of the SHP-1000+. Strong attenuation below 1 GHz and low insertion loss in the 1–2 GHz band are both visible.

3.3.2 Noise Figure and the Friis Formula

The main question for any receiver is how faint a signal it can detect. The answer depends on how much noise the receiver itself adds on top of the sky noise. This is described by the *noise figure* (NF) of the receiver, defined as the ratio of the input signal-to-noise ratio (SNR) to the output SNR, expressed in decibels:

$$\mathrm{NF} = 10\log_{10}(F)\ \mathrm{dB}, \qquad F = \frac{\mathrm{SNR}_{\mathrm{in}}}{\mathrm{SNR}_{\mathrm{out}}} \tag{3.13}$$

A noise figure of 0 dB means the receiver adds no noise at all, which is impossible in practice. Lower values are better. The equivalent noise temperature is:

$$T_e = T_0\,(F-1), \qquad T_0 = 290\ \mathrm{K} \tag{3.14}$$

When several components are connected in series, the total noise factor is given by the *Friis formula* [14]:

$$\boxed{F_{\mathrm{total}} = F_1 + \frac{F_2-1}{G_1} + \frac{F_3-1}{G_1G_2} + \cdots + \frac{F_n-1}{G_1G_2\cdots G_{n-1}}} \tag{3.15}$$

where F_i and G_i are the noise factor and available power gain of stage i, both expressed in linear units. The key implication of this equation is that the first stage dominates the overall noise performance. The noise contribution of each later stage is reduced by the gain of the stages

before it. If the first LNA has a gain of $G_1 = 8913$ (i.e. 39.5 dB), then the noise contribution of the second stage is divided by 8913 and is therefore negligible for any reasonable component.

For our LNA ($F_1 = 10^{0.04/10} \approx 1.009$, $G_1 \approx 8913$ linear):

$$\frac{F_2 - 1}{G_1} \approx \frac{1.585 - 1}{8913} \approx 6.6 \times 10^{-5} \tag{3.16}$$

The contribution of the BPF to the total noise factor is therefore less than 0.01%, and is completely dominated by the first LNA. This is why placing a high-gain, low-noise amplifier as the first component after the antenna is so important [14].

3.3.3 Front-End Configurations Evaluated

Three orderings of the four front-end components were assembled and tested using a Vector Network Analyser (VNA). The VNA measures how signals propagate through the chain as a function of frequency, particularly the gain (S_{21}) and return loss (S_{11}). Comparing these configurations helped verify that the measured behaviour agreed with the expectations from the Friis formula.

Configuration 1 - Optimal Low-Noise (Selected)

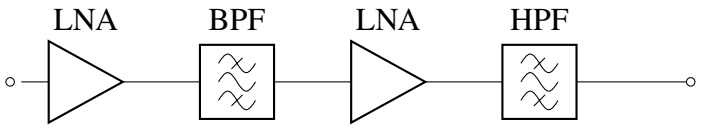

In this configuration, the first LNA is placed directly after the antenna. By the Friis formula, this minimises the system noise figure. The BPF suppresses interference between the two gain stages, and the HPF removes low-frequency interference before the backend. This configuration was used for all astronomical observations in this thesis.

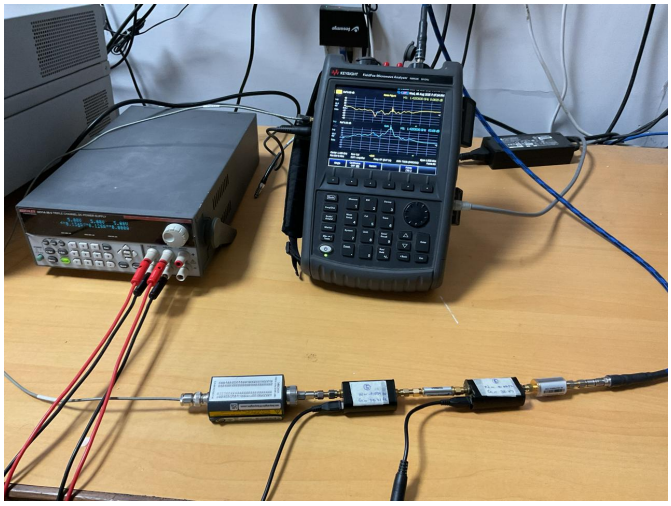

Figure 3.14: VNA measurement of Configuration 1: LNA → BPF → LNA → HPF. The LNA-first ordering gives the lowest system noise figure (Friis formula, Eq. 3.15) and was selected for our system.

Configuration 2 — Filter-First (RFI Protection)

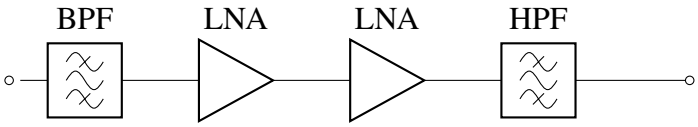

Here, the BPF is placed before the first LNA. This improves protection against strong out-of-band interference that could saturate the amplifier. However, the filter insertion loss of about 2 dB now appears before the first gain stage, increasing the system noise figure by about 2 dB. This penalty is significant for radio astronomy, but the configuration can still be useful in high-RFI environments.

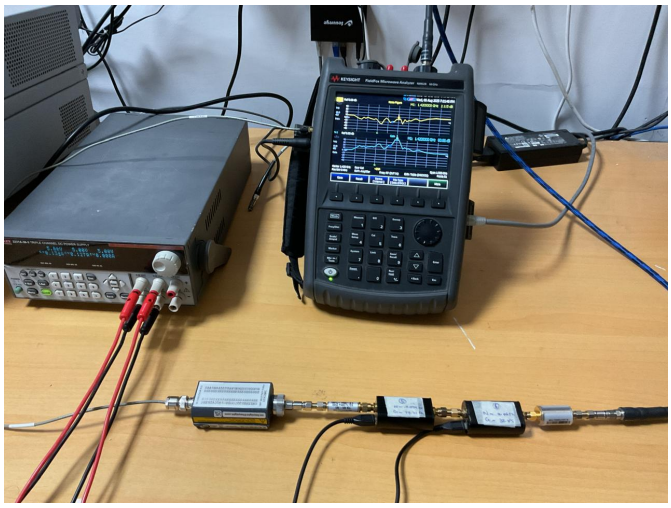

Figure 3.15: VNA measurement of Configuration 2: BPF → LNA → LNA → HPF. Placing the filter first protects against strong RFI but raises the noise floor by ≈ 2 dB.

Configuration 3 — Attenuator-Stabilised

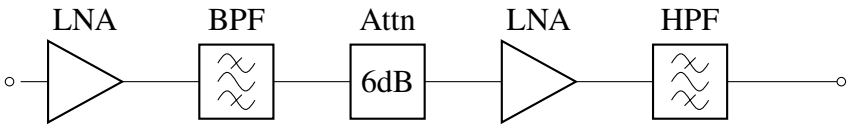

In this configuration, a fixed 6dB attenuator is inserted between the two gain stages. It improves impedance matching between the BPF output and the second LNA input, helping prevent nonlinear compression during transient RFI events. This improves stability during long observations. Since the attenuator is placed after the first LNA, its effect on the overall system noise figure is negligible by the Friis formula. This configuration was tested once for stability, but Configuration 1 was retained for the final observations.

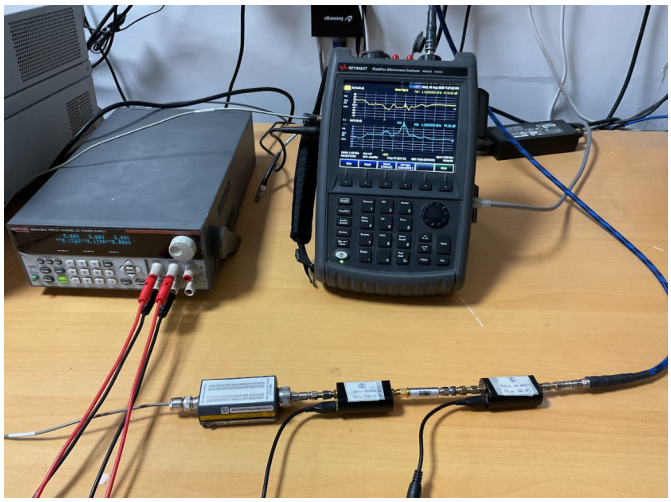

Figure 3.16: VNA measurement of Configuration 3: LNA → BPF → Attn → LNA → HPF. The attenuator improves stability for long observations. System noise figure impact is negligible because the attenuator sits after the first LNA.

3.3.4 Front-End Configuration Selection

The receiver front-end strongly affects the sensitivity of the 1.42GHz hydrogen-line system. To optimise the signal chain, three configurations were tested by rearranging the same components: low-noise amplifiers (LNAs), band-pass filters, and an attenuator. The aim was to maximise gain and minimise noise within the 1350–1450MHz band while maintaining a stable passband.

The gain response of each configuration was measured using a vector network analyser (VNA) over 700–2200MHz. The comparison is shown in Figure3.17.

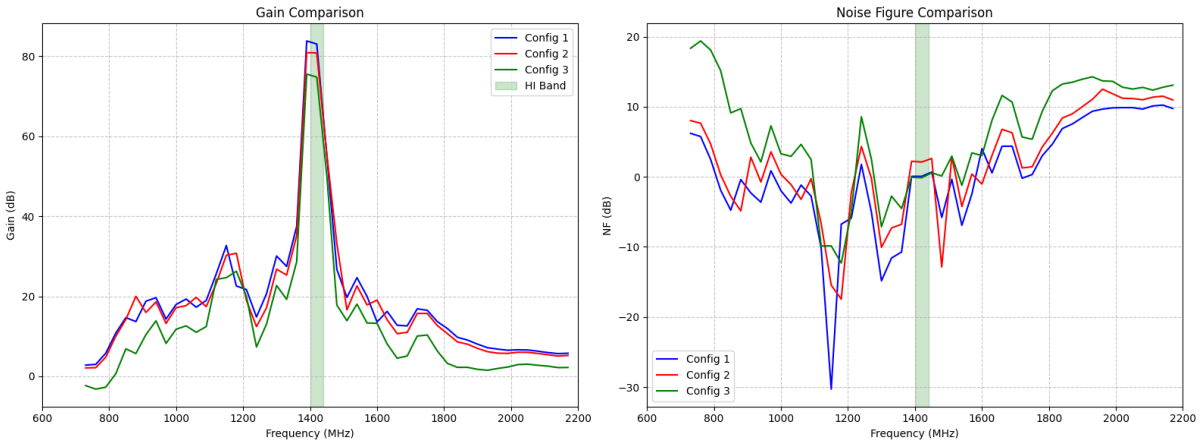

Figure 3.17: VNA-measured gain comparison of the three front-end configurations.

Configuration 1 shows the highest and most stable gain within the hydrogen-line band. Because the first-stage LNA is placed directly at the input, its low noise figure dominates the overall system performance. From the Friis relation (Eq. 3.15), the noise contribution of later stages is strongly reduced by the gain of the first stage.

In Configuration 2, placing the band-pass filter before the first LNA introduces insertion loss ahead of amplification, which increases the effective noise figure and reduces sensitivity. Configuration 3 includes an inter-stage attenuator, which improves impedance matching but reduces the overall gain.

The selected configuration provides a total gain of $\sim 80\,\mathrm{dB}$, sufficient to bring the weak astronomical signal into the input range of the SDR or RFSoC ADC. At the same time, the system noise remains dominated by the first LNA, preserving near-optimal sensitivity.

For these reasons, Configuration 1 was used for all observations. It maximises early-stage gain, suppresses downstream noise contributions, and maintains a stable passband for long-duration interferometric measurements.

Chapter 4

Digital Backend

The horn antennas and front-end electronics (Chapter 3) collect faint radio emission, amplify it by 80 dB, and deliver a filtered analogue signal centred at 1.42 GHz. At this stage, the signal is still a high-frequency voltage and cannot be processed directly. The digital backend converts it into sampled data, performs channelisation and cross-correlation, and records the visibilities for analysis.

Developing the backend was the most challenging part of the project and required significant changes from the initial design. Two approaches were implemented and tested. Early experiments with low-cost RTL-SDR dongles showed that they do not provide the phase coherence required for interferometry. The system was then moved to a USRP B210 software-defined radio, which provided stable phase-coherent performance and enabled clear solar fringe detection. All astronomical results in Chapter 5 are based on this configuration. The system was further extended at the bench level to four elements using two synchronised USRP B210 units with a shared reference clock and PPS signal; this configuration is described in Section 4.2.3 and validated in Section 5.2.4. This chapter presents both configurations and summarises their performance and the main lessons learned.

4.1 Backend 1: RTL-SDR

4.1.1 Motivation

The natural first choice for a low-cost educational interferometer is the RTL-SDR dongle. Originally designed for digital television reception, the RTL2832U chipset inside these compact USB devices can be repurposed as a general-purpose IQ receiver covering roughly 25–1750 MHz [9]. At approximately 1000 per unit, they are by far the most affordable SDR hardware available, and several groups have successfully used them for hydrogen-line spectroscopy and galactic rotation curve measurements [9]. The concept was straightforward. One dongle was assigned to each antenna, and the outputs were correlated in GNU Radio to form a low-cost two-element

interferometer. At first glance, the approach appeared simple.

4.1.2 The Phase Coherence Requirement

Interferometry depends on measuring the phase difference between signals received at two antennas, which is meaningful only if both receivers share a common time reference. Each RTL-SDR dongle uses an independent temperature-compensated crystal oscillator (TCXO) without external locking. As a result, two units drift relative to one another, and the measured phase becomes dominated by oscillator instability rather than the sky signal.

A partial workaround is clock daisy-chaining, in which one dongle acts as the master and its 28.8 MHz TCXO is routed to a second unit, replacing the second dongle's internal oscillator. In this arrangement, both devices share a common reference frequency.

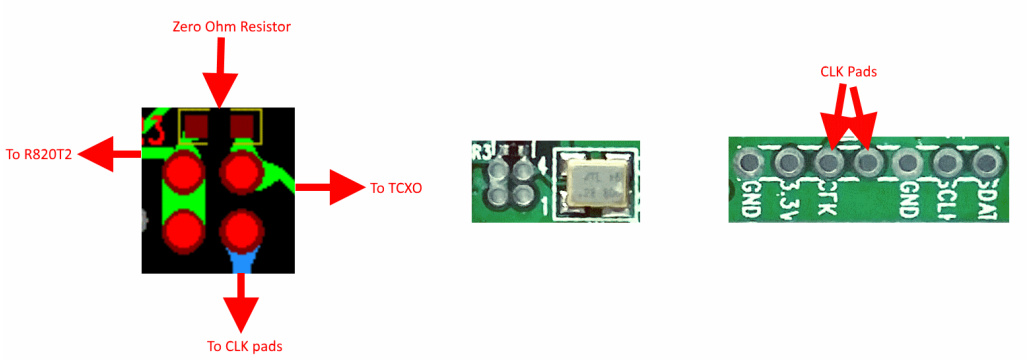

(a) CLK pad on the RTL-SDR Blog V3 PCB. The zero-ohm resistor connects the internal TCXO to the output; removing it allows an external reference to be injected.

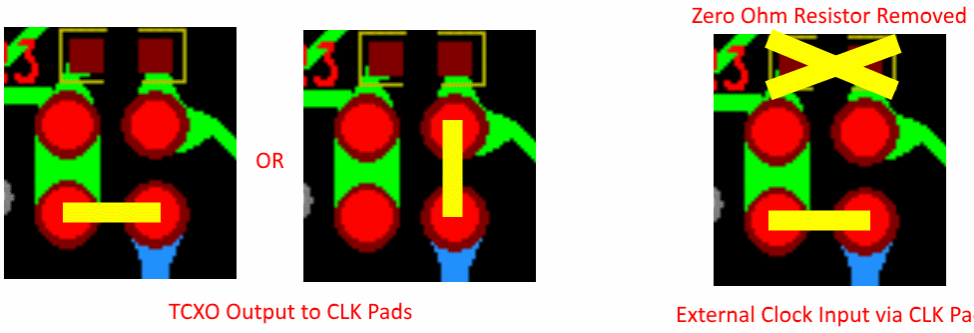

(b) Clock configurations: internal TCXO driving the CLK pad (left), and external reference injection after removing the resistor (right).

Figure 4.1: Clock routing on the RTL-SDR Blog V3. A shared reference frequency can be provided in hardware, but as the bench tests show, this does not guarantee phase coherence [15].

4.1.3 Bench Test Setup

To test coherence, a controlled bench experiment was carried out before antenna observations. A 1420.405 MHz tone was split using a two-way power divider and fed identically to both dongles. Because the inputs were the same, the cross-correlation phase was expected to remain constant. Any variation therefore reflects instrumental effects.

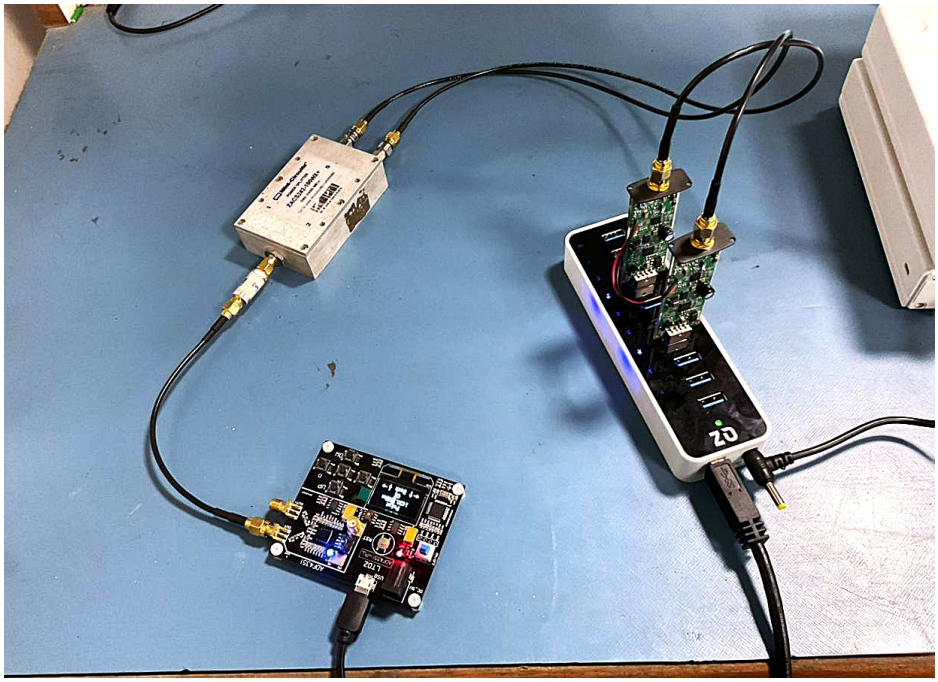

Figure 4.2: Bench test setup for the RTL-SDR phase stability measurement. Two RTL-SDR dongles receive the same split signal via a two-way RF power splitter. Both are connected to the host computer through a powered USB hub. Any variation in the measured phase difference between the two channels is attributable entirely to the receivers.

4.1.4 GNU Radio Flowgraph

The cross-correlation was implemented in GNU Radio using the FX architecture. Each dongle was accessed via a separate RTL-SDR source block: the master (rtl=0) used its internal clock, and the slave (rtl=1) used the external clock input to receive the daisy-chained reference. Each channel was passed through an 8192-point Blackman-Harris windowed FFT. The two spectral streams were then multiplied conjugately channel by channel, and the resulting cross-power spectrum was decomposed into magnitude and phase using a Complex-to-Mag/Phase block. Both outputs were displayed in real time using QT GUI Time Sinks.

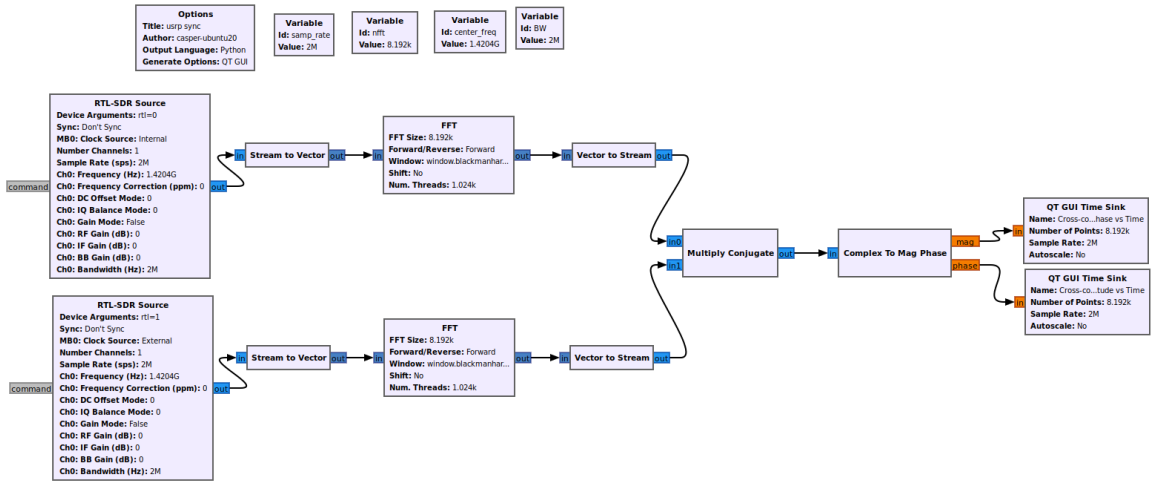

Figure 4.3: GNU Radio flowgraph for the RTL-SDR phase coherence test. Both dongles stream into separate FFT paths, after which a Multiply Conjugate block forms the cross-power spectrum. The magnitude and phase of the cross-product are displayed in real time.

4.1.5 Result and Interpretation

The result of the bench test is described in Section 5.1.1 of the results chapter. In short, the cross-correlation phase showed no stable value and wandered continuously across the full $[-\pi, +\pi]$ range throughout the observation, confirming that the daisy-chain does not produce phase-coherent outputs.

The root cause is a fundamental property of the hardware. The daisy-chain provides a common reference *frequency* to both dongles, but not a common *phase*. Inside each dongle, a phase-locked loop generates the local oscillator from the 28.8 MHz reference. A PLL locks its output frequency to the reference, but the phase at which it settles at power-up is random, and its accumulated phase noise evolves independently of any other PLL. Two PLLs locked to the same reference frequency run at the same rate, but their relative phase is arbitrary at startup and may drift with time. This is a hardware limitation, not a software issue. Reliable interferometry therefore requires the two receive channels to share not only a reference frequency but also a *single local oscillator*. This requirement motivated the use of the USRP B210, in which both channels are generated by a common AD9361 chip and share a single internal oscillator by design.

4.2 Backend 2: USRP B210 Software-Defined Radio

4.2.1 Hardware Overview

The Ettus Research USRP B210 is a software-defined radio platform that operates over a continuous frequency range from 70 MHz to 6GHz. Designed for scientific and engineering

applications, it provides two receive and two transmit channels based on a single Analog Devices AD9361 RF transceiver. The device connects to a host computer through USB3.0, with all signal processing and analysis carried out in software.

The architecture of the AD9361 chip is what makes this device suitable for interferometry. Both receive channels share a single internal local oscillator (LO), a single set of phase-locked loop (PLL) synthesisers, and a single 40MHz temperature-compensated crystal oscillator (TCXO) reference. As a result, the two channels are inherently phase-coherent from power-up. They observe the same carrier frequency using the same oscillator, with no independent drift between them. Achieving this property with two separate RTL-SDR receivers required hardware modifications that ultimately failed. In the B210, it is a built-in feature. Table4.1 summarises the main specifications relevant to this project.

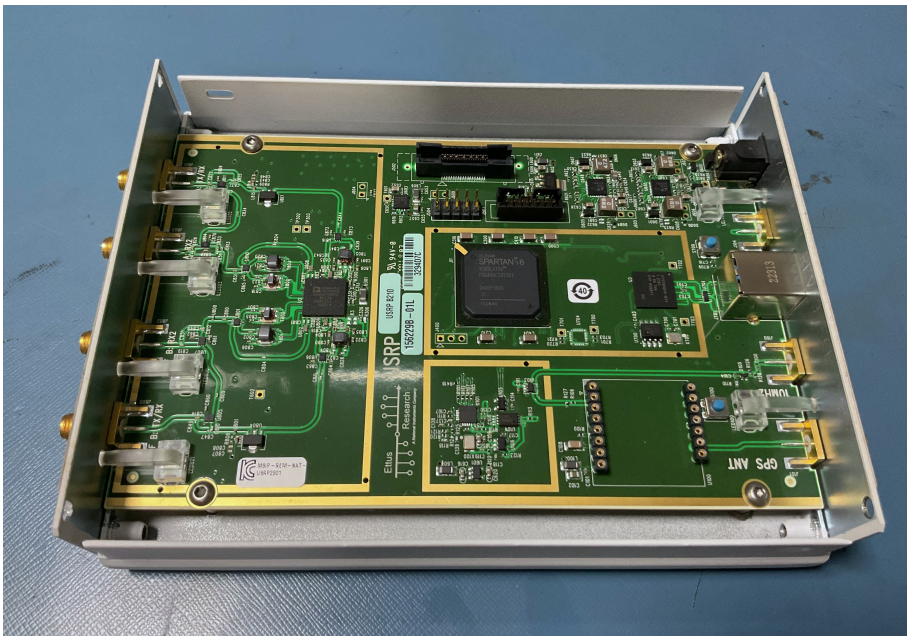

Figure 4.4: The Ettus Research USRP B210. Both receive channels share a single AD9361 transceiver chip and internal clock, providing inherent two-channel phase coherence. The REF IN and PPS IN SMA connectors allow locking to an external frequency reference and pulse-per-second timing signal for multi-device synchronisation.

Table 4.1: Key specifications of the USRP B210 relevant to this project.

Parameter	Value
Frequency range	70 MHz to 6 GHz
Number of receive channels	2 (fully coherent, single AD9361)
ADC resolution	12 bits
Maximum sample rate (complex IQ)	61.44 Msps
USB interface	USB 3.0
External reference input (REF IN)	10 MHz SMA connector
Timing input (PPS IN)	Pulse-per-second SMA connector
Software driver	UHD (USRP Hardware Driver, open source)

4.2.2 Single USRP B210: Two-Element Phase Coherence Verification

Before any outdoor astronomical observations were made, the internal phase coherence of the two B210 channels was tested in a controlled bench setup. This step was necessary to experimentally confirm, rather than relying only on the datasheet, that the device behaves as a genuinely phase-coherent receiver pair.

Bench Test Setup

A continuous wave (CW) tone at 1420.405MHz was generated by a signal generator and split through a two-way RF power divider. One output was connected to the RFA RX2 port (Channel 0) and the other to the RF B RX2 port (Channel 1) of a single USRP B210. Since both channels received exactly the same signal, the measured cross-correlation phase between them should remain constant over time. Any variation is purely instrumental.

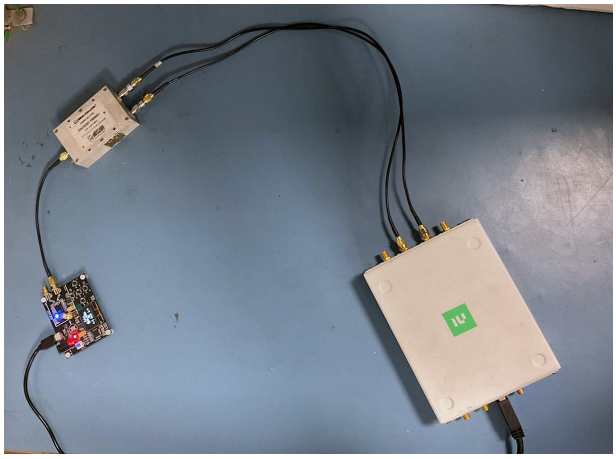

(a) Single USRP B210 coherence test setup with signal generator and power divider feeding both RX2 inputs.

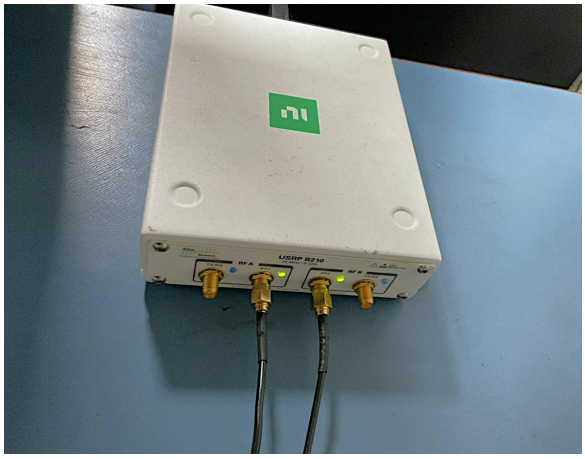

(b) RF connections to RX2 ports (Channel 0 and Channel 1) receiving identical input signals.

Figure 4.5: Single USRP B210 phase coherence test: (a) setup overview; (b) RX2 input connections.

FX Correlator Architecture

All GNU Radio correlator flowgraphs in this project implement the FX architecture: the Fourier transform is applied first (F step), followed by channel-by-channel cross-multiplication (X step). This ordering is both mathematically correct and computationally efficient.

In the F step, $N = 8192$ consecutive complex IQ samples are collected from each channel and transformed by a Blackman–Harris windowed FFT, giving 8192 spectral channels with a frequency resolution of:

$$\Delta\nu = \frac{f_s}{N} = \frac{2{,}000{,}000}{8192} \approx 244 \text{ Hz} \tag{4.1}$$

The Blackman–Harris window suppresses spectral sidelobes to below -92 dB, preventing strong RFI in one channel from contaminating neighbouring channels.

In the X step, the cross-power spectrum is computed channel by channel:

$$S_{12}[k] = \tilde{V}_1[k] \cdot \tilde{V}_2^*[k], \quad k = 0, 1, \ldots, N-1 \tag{4.2}$$

The amplitude $|S_{12}[k]|$ is the correlated power and the phase $\arg(S_{12}[k])$ is the interferometric visibility phase. A moving average over successive cross-power spectra integrates the sky signal while averaging down uncorrelated receiver noise.

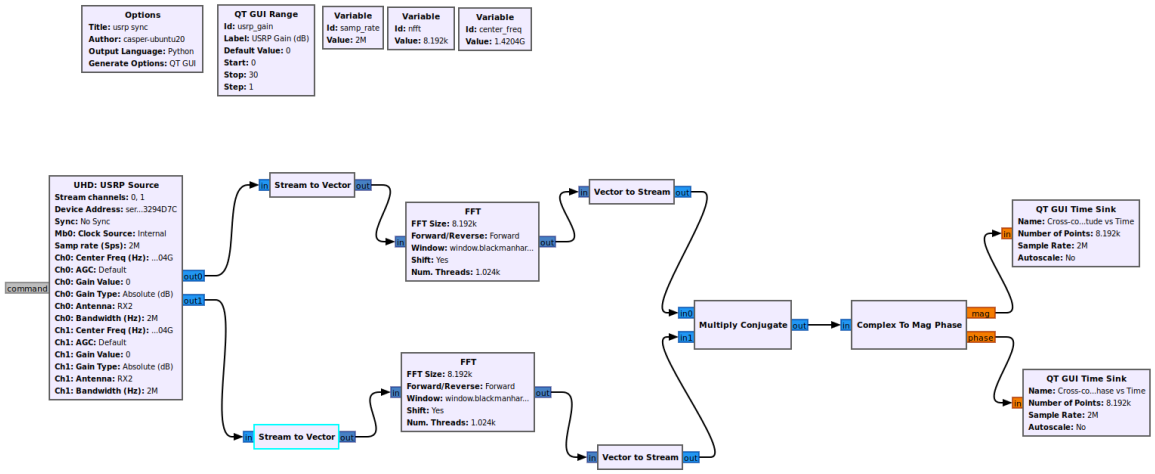

Figure 4.6: GNU Radio Companion FX correlator flowgraph for single USRP B210 phase coherence verification. Both channels are processed through parallel FFT chains, and the cross-power spectrum is computed using a Multiply Conjugate block. The resulting visibility magnitude and phase are displayed in real time.

The test was implemented using an FX correlator flowgraph in GNU Radio Companion. Both channels are streamed simultaneously from a single UHD source block, ensuring a shared local oscillator. Each channel passes through a Stream-to-Vector block, an 8192-point Blackman–Harris windowed FFT, and a Vector-to-Stream block. The two spectral streams enter a Multiply Conjugate block, which computes the cross-power spectrum $S_{12}[k] = \tilde{V}_1[k] \cdot \tilde{V}_2^*[k]$ for all channels simultaneously. A Complex-to-Mag/Phase block extracts the visibility magnitude and phase, displayed in real time using QT GUI Time Sinks. The observed phase stability is presented in Chapter 5.

4.2.3 Extending to Four Elements: Two Synchronised USRP B210s

Why Two USRPs Cannot Simply Run Together

With a single USRP B210 operating reliably as a two-element interferometer, the next step was to extend the system to four antenna elements using a second USRP B210. This increased the number of simultaneous baselines from one to six, significantly improving UV-plane coverage.

The limitation is fundamental. Each USRP B210 has its own independent 40MHz TCXO oscillator. When two units are powered and operated without an external reference, each follows its own clock independently. As a result, the relative phase between their LOs is random at startup and introduces a time-varying phase due to frequency offset between the two independent oscillators, making long-term interferometric cross-correlation between the two devices unreliable. This is the same issue that caused the RTL-SDR daisy-chain approach to fail

in Section4.1, although here it appears at the level of two USRP devices.

To confirm this experimentally, the two USRP units were connected to a split 1420.405 MHz tone with both devices running on their internal clocks and no external synchronisation applied. The GNU Radio flowgraph used for this test is shown in Figure 4.7.

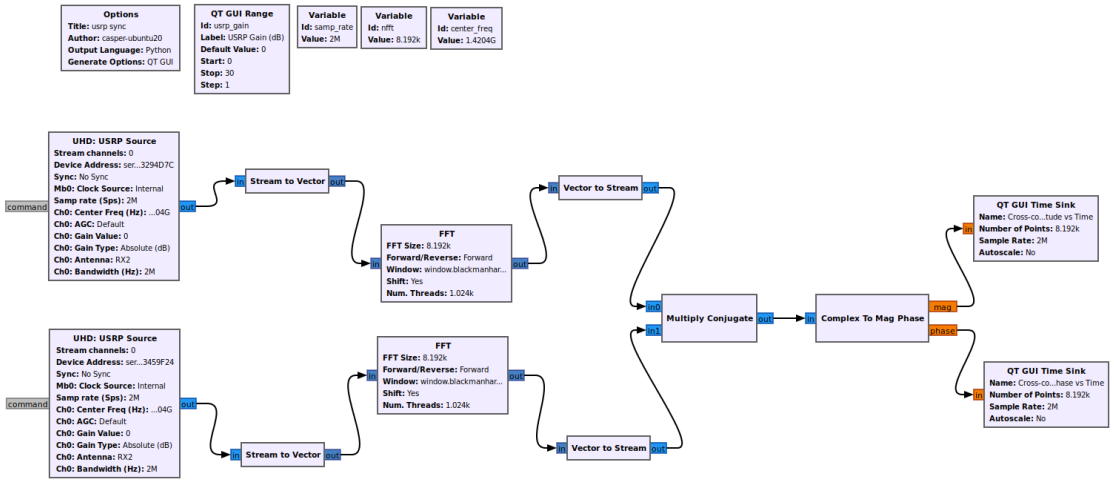

Figure 4.7: GNU Radio FX correlator flowgraph for the two-USRP incoherence baseline test. Both USRP sources operate on their internal clocks with no synchronisation (Clock Source: Internal, Sync: No Sync). Each channel passes through an 8192-point Blackman–Harris FFT before cross-multiplication. The cross-power magnitude and phase are displayed in real time.

The result, shown in Chapter 5, exhibits a phase trace that sweeps across the full $[-\pi, +\pi]$ range, indicating that no stable phase relationship exists between the two devices. This baseline test motivated the synchronisation approach described below.

External Reference and PPS

The USRP B210 supports inter-device synchronisation through an external frequency reference and a pulse-per-second (PPS) signal, both provided through rear-panel SMA connectors.

REF IN (10 MHz reference). A 10 MHz signal applied to the REF IN port replaces the internal oscillator with an external reference. When both devices share the same source, they operate at a common frequency, removing relative drift. A clean sine wave is preferred for stable PLL locking.

PPS IN (timing signal). A 1 Hz square wave applied to the PPS IN port provides a timing reference for aligning device timestamps. The UHD driver uses the rising edge to synchronise

the internal time counters. Without PPS, the devices may begin sampling at different times even if they are frequency-locked.

For reliable operation, the PPS signal should be TTL-level (0–3.3 V), with a pulse width of a few tens of milliseconds, and should be derived from the same timebase as the 10 MHz reference.

Signal Source: Moku:Pro

Both the 10 MHz reference and PPS signals were generated using a Liquid Instruments Moku:Pro in waveform generator mode. One output produced a 10 MHz sine wave (approximately 1.4 Vpp into 50 Ω), while the other produced a 1 Hz square wave with a +1.65 V offset, giving a 0–3.3 V PPS signal. The duty cycle was set to ~ 25% (pulse width ~ 20–25 ms).

Because both outputs share a common internal timebase, the PPS signal remains phase-consistent with the 10 MHz reference, providing a coherent timing and frequency reference for both USRP devices.

Although the Moku:Pro was used here for convenience, any source capable of providing a stable 10 MHz reference and a synchronised PPS signal, such as a GPS-disciplined oscillator, would be suitable.

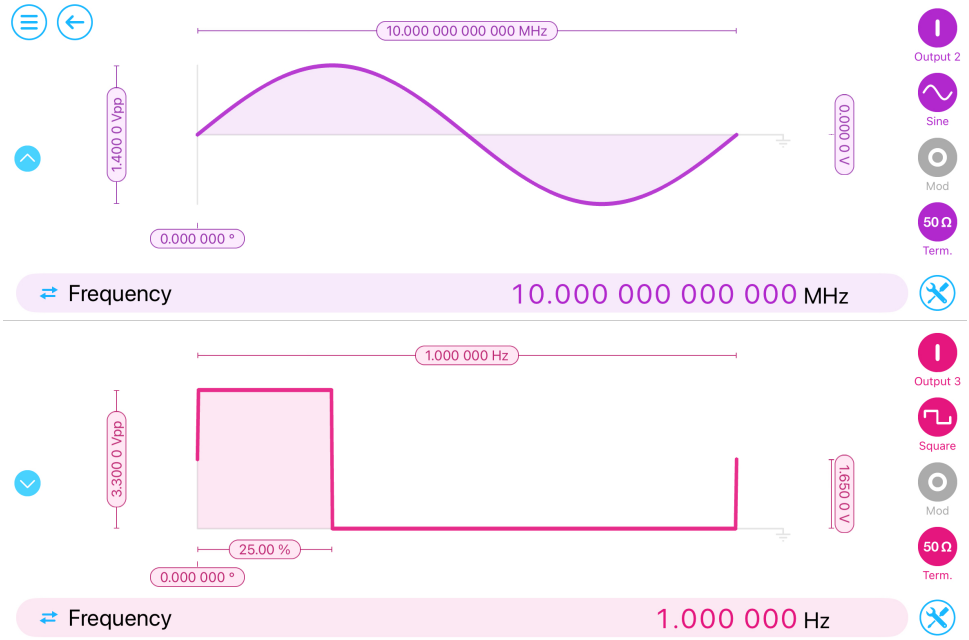

Figure 4.8: Moku:Pro waveform generator configuration used for synchronisation. The upper panel shows the 10 MHz reference signal, while the lower panel shows the 1 Hz PPS signal with a DC offset producing a 0–3.3 V waveform. Both outputs share a common internal timebase.

Hardware Setup: Baseline 1–3 (Two-Channel Sync Test)

The first inter-device synchronisation test used one channel from each USRP: Channel0 of USRP1 (serial `3459F24`) and Channel0 of USRP2 (serial `3294D7C`). A single 1420.405MHz tone was split through a two-way power divider and fed to the RX2 port of each device. The 10MHz reference from the Moku:Pro was distributed to the REF IN port of both USRPs through a power splitter, and the PPS output was similarly distributed to both PPS IN ports.

GNU Radio Time-Domain Sync Flowgraph

The time-domain synchronisation test extends the single-device verification to two independent USRP B210 units. Each device is accessed through a separate UHD source block, identified by its serial number, with both Clock Source and Time Source configured to `External` to enable a shared reference and PPS alignment.

The complex baseband signals from the two USRPs are processed in parallel. A short delay block provides coarse alignment, followed by a moving average filter to reduce noise. The signals are then separated into real, imaginary, magnitude, and phase components using standard GNU Radio blocks. These quantities are displayed using QT GUI Time Sinks and recorded for offline analysis. This allows direct comparison of amplitude and phase between the two devices, providing an initial time-domain assessment of inter-device synchronisation before frequency-domain correlation.

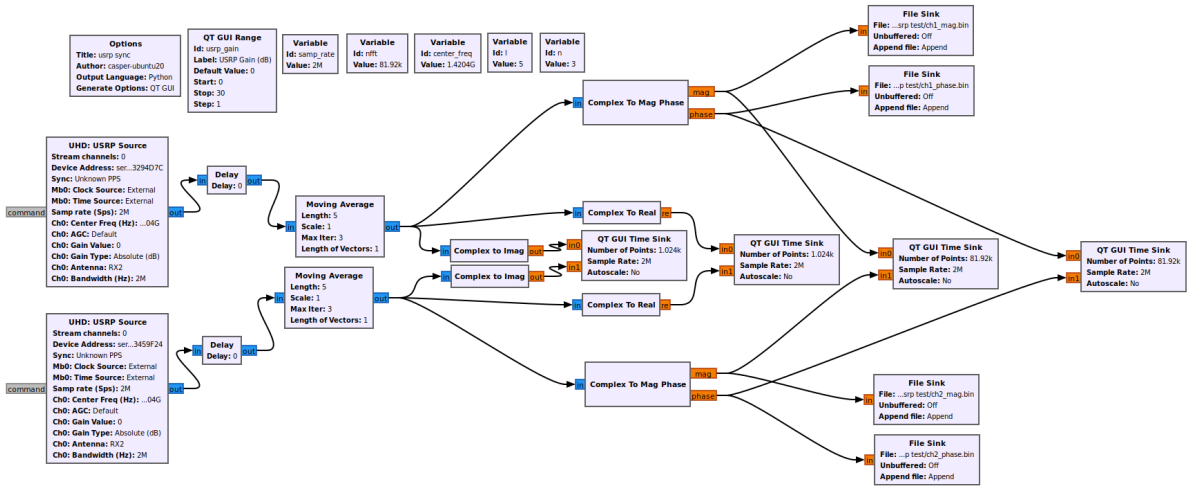

Figure 4.9: GNU Radio flowgraph for two-USRP time-domain synchronisation testing. Each USRP is accessed through a separate UHD source block with external clocking and PPS. The signals are processed in parallel, and their real, imaginary, magnitude, and phase components are displayed and recorded.

4.2.4 Synchronisation: Issues Encountered and Solutions

Achieving phase-coherent operation between two USRP B210 units required more than supplying a shared 10 MHz reference and PPS signal. Although these provide a common frequency and timing reference, repeatable phase alignment also depends on how device timing is controlled in software.

Both devices were configured to use external clock and time sources, and correct locking was verified through the UHD driver. Even so, the relative phase varied between runs when the default GNU Radio Companion (GRC) setup was used.

Limitations of GNU Radio Companion

GRC does not provide enough control over multi-device timing. In the generated Python code, timing initialisation is handled internally, with no guarantee that both devices align to the same PPS edge. As a result, the phase offset remained unpredictable between runs, even with external clocking.

To address this, the flowgraph was exported to Python and modified to explicitly control timing through the UHD API, while leaving the signal-processing chain unchanged.

PPS-Based Time Alignment

Providing a PPS signal alone does not guarantee that both devices initialise on the same PPS edge. If timing commands are issued sequentially, the devices may latch onto different PPS events, producing arbitrary phase offsets.

In addition, the default `set_time_unknown_pps()` calls generated by GRC occasionally caused UHD timeouts because they were executed without waiting for a valid PPS transition.

These calls were removed and replaced with an explicit sequence: waiting for a PPS edge using `get_time_last_pps()`, issuing `set_time_next_pps(0.0)` to both devices, and scheduling a common start time with `set_start_time()`. This ensured that both devices reset and began sampling on the same PPS edge.

Thread Timing and Residual Offset

To minimise delay, `set_time_next_pps()` was issued to both devices in separate threads. Because of the Python GIL, a small offset ($\sim 10^2$ μs) remains between the commands.

This produces a fixed phase offset but does not cause phase drift. In practice, it does not limit interferometric measurements, since it can be removed through standard calibration. Once aligned, both devices share a common time reference and show stable phase behaviour.

Buffer Overflows and Data Throughput

During the initial four-channel tests, the terminal intermittently displayed buffer overflow warnings as repeated `O` symbols. These warnings indicate dropped data buffers and cause discontinuities in the recorded phase.

The primary cause was insufficient USB bandwidth, as one of the devices was initially connected via USB 2.0. Moving both USRPs to USB 3.0 ports resolved most of the issue.

A second contributing factor was the use of a large FFT size, which increased the processing load on the GNU Radio scheduler. Reducing `nfft` to 4096 and increasing the receive buffer depth by adding `num_recv_frames=256` to the UHD device arguments eliminated the remaining overflows and resulted in stable continuous operation.

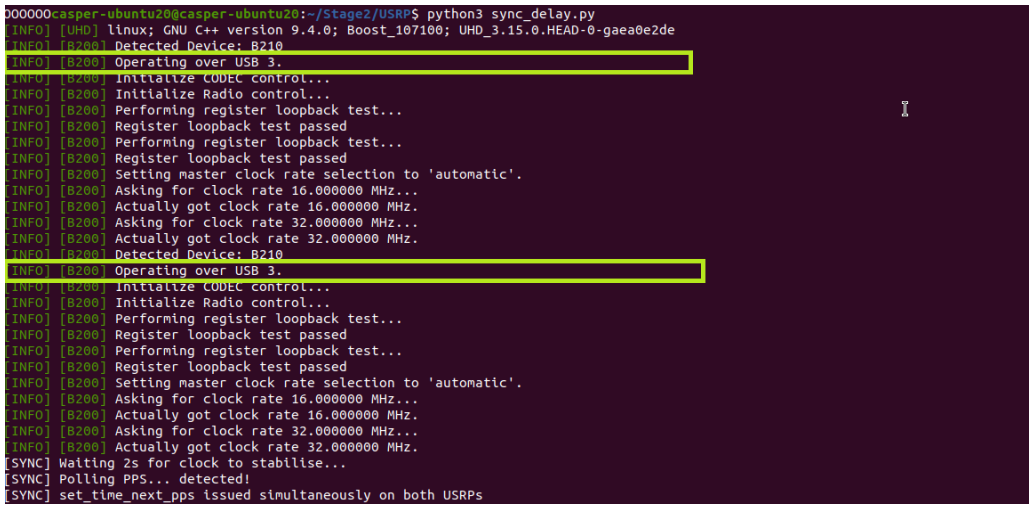

Figure 4.10: Terminal output confirming both USRP B210 units are operating over USB 3.0.



(a) Overall bench setup for the two-USRP synchronisation test. Both USRP B210 units receive the same split tone. REF IN and PPS IN signals from the Moku:Pro are distributed to both units via power splitters.



(b) Signal injection in USRP inputs: the 1420.405 MHz tone from the signal generator is fed through a two-way RF power splitter. One output goes to Ch0 of USRP 1 (RF A RX2) and the other to Ch0 of USRP 2 (RF A RX2).

Figure 4.11: Bench setup overview for the two-USRP synchronisation test. (a) Overall system layout. (b) RF signal injection path.



Figure 4.12: Rear panel of one USRP B210 showing the USB 3.0 data connection, the 10 MHz REF IN input, and the PPS IN input. The REF IN and adjacent GPS port are physically identical SMA connectors; connecting to the wrong port is a common source of error during synchronisation setup.

4.2.5 FFT Flowgraph: Baseline 1–3 (Two-USRP, Single Baseline)

For the frequency-domain synchronisation verification of Baseline 1–3, a second flowgraph was built: the FX correlator. This flowgraph implements exactly the same architecture used for the two-element astronomical observations, making it a direct functional test of the actual instrument pipeline.

Both USRP B210 units stream Channel 0 at 2 Msps centred on 1420.405 MHz with Clock

Source and Time Source set to `External`. Each channel passes through Stream-to-Vector, an 8192-point Blackman-Harris windowed FFT, and Vector-to-Stream. The two FFT output streams enter a Multiply Conjugate block, which computes $S_{12}[k] = \tilde{V}_1[k] \cdot \tilde{V}_2^*[k]$ for all 8192 channels simultaneously. A Complex to Mag Phase block splits the result into magnitude and phase, and two QT GUI Time Sink blocks display both quantities in real time. The flowgraph was then exported as a Python script and the synchronisation modifications were applied.

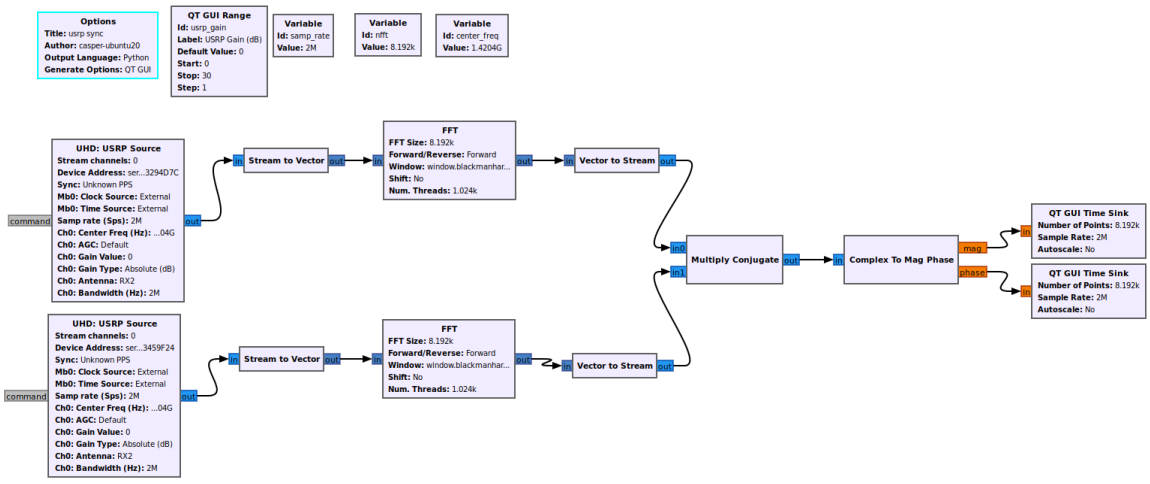

Figure 4.13: GNU Radio Companion flowgraph for frequency-domain synchronisation verification of Baseline 1–3. Both USRP B210 units stream Channel 0 with external clock and time sources. Signal flow: Stream to Vector → FFT (8192-point, Blackman-Harris) → Vector to Stream → Multiply Conjugate → Complex to Mag Phase → QT GUI Time Sink. The cross-correlation output from this flowgraph is presented in Chapter 5

4.2.6 Four-Element Six-Baseline Configuration

Channel Mapping and Baseline Definitions

With both USRP B210 units providing two receive channels each, the system has four channels in total. Following the labelling used throughout Chapter 5, these are designated A1 and A2 (USRP 1, serial `3459F24`, channels 0 and 1) and B1 and B2 (USRP 2, serial `3294D7C`, channels 0 and 1). The six cross-correlation baselines formed from these channels are listed in Table 4.2.

Baselines 1–2 and 3–4 are formed entirely within a single AD9361 chip and are inherently phase-coherent regardless of the external synchronisation state. They therefore serve as a real-time coherence reference during every four-element run. Baselines 1–3, 1–4, 2–3, and 2–4 cross between the two devices and require the external synchronisation to achieve coherent operation.

Table 4.2: Channel mapping and baseline definitions for the four-element system. Internal baselines are inherently phase-coherent; cross-device baselines require external synchronisation.

Baseline	Channels	Type	Coherence
1–2	A1 × A2*	Internal (USRP 1)	Inherent
1–3	A1 × B1*	Cross-device	External sync required
1–4	A1 × B2*	Cross-device	External sync required
2–3	A2 × B1*	Cross-device	External sync required
2–4	A2 × B2*	Cross-device	External sync required
3–4	B1 × B2*	Internal (USRP 2)	Inherent

*Complex conjugate of the second channel.

Hardware Setup for Four Elements

For the four-element test, a 1420.405 MHz tone from the signal generator was first amplified by an LNA and then fed through a four-way RF power splitter. One output was connected to each of the four receive ports: RF A RX2 and RF B RX2 on USRP 1, and RF A RX2 and RF B RX2 on USRP 2. REF IN and PPS IN connections to both units were identical to the Baseline 1–3 setup described above.

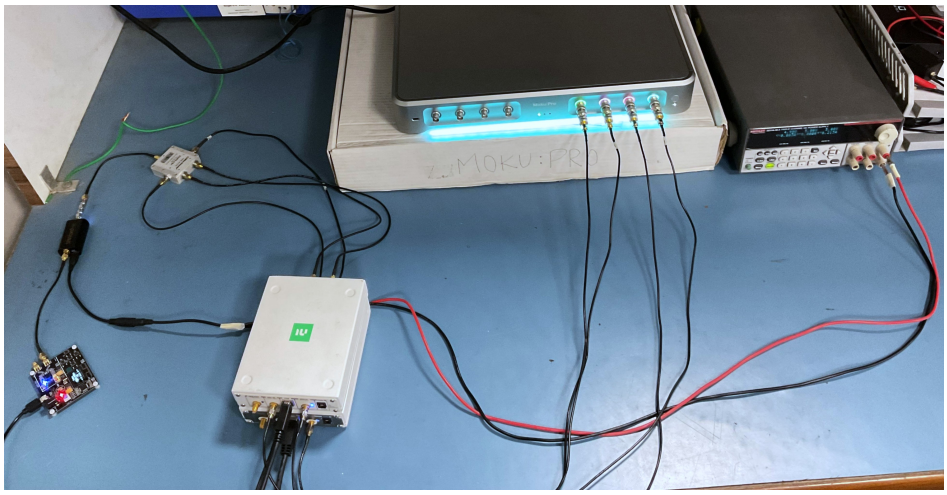

Figure 4.14: Overall bench setup for the four-element interferometer test, showing the USRP B210 units, signal generator, RF 4-way splitter, and reference timing distribution via the Moku:Pro.



Figure 4.15: Close-up of the four USRP receive ports used in the four-element test. The RF A RX2 and RF B RX2 inputs on both USRP 1 and USRP 2 are each connected to one output of the four-way RF splitter, ensuring identical signal injection across all channels.

Four-Element GNU Radio Flowgraph

The four-element correlator flowgraph extends the two-channel FX pipeline to four channels and computes all six cross-products. Each of the four channels has its own Stream-to-Vector and FFT path. After the FFT, six Multiply Conjugate blocks compute the cross-power spectrum for each baseline, and six Complex to Mag/Phase blocks extract the visibility magnitude and phase. Two QT GUI Time Sink blocks display all six magnitude traces and all six phase traces simultaneously.

Both USRP source blocks stream both channels by setting `channels=[0, 1]` in the stream arguments. The same Python synchronisation block used for the two-channel test was retained without modification: the timing procedure is identical regardless of how many channels are streamed from each device.

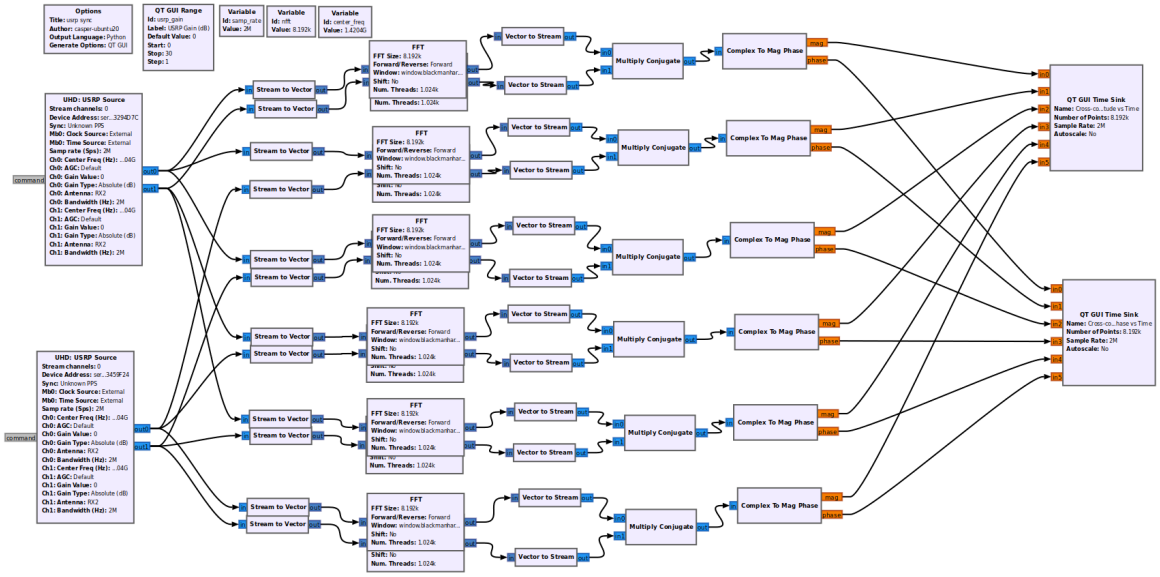

Figure 4.16: GNU Radio four-element FX correlator flowgraph. Four independent FFT paths (two channels from each USRP) feed into six Multiply Conjugate blocks computing all six baselines simultaneously. The magnitude and phase of all six baselines are displayed in real time and written to file sinks. The results from this flowgraph are presented in Chapter 5.

4.2.7 Current Status and Limitations of the Four-Element System

The four-element system demonstrated phase-coherent operation across all six baselines in controlled bench tests. A shared 10 MHz reference provided a common frequency, while the PPS-aligned timed-start synchronised the CORDIC accumulators so that both devices began sampling simultaneously.

A fixed LO phase offset remained after each power cycle. When two AD9361 chips lock independently to the same reference, their PLLs settle to a random but stable relative phase. This offset remains constant within a session but changes between sessions. It can be measured using a calibration signal at the start of each observation and corrected in post-processing, following standard interferometric practice.

The system has not yet produced astronomical fringe data. The next step is to extend the setup to on-sky observations with appropriate calibration.

4.3 Outdoor Solar Drift Scan Observations

After bench validation of the two-element correlator, the system was deployed on the IIT Indore rooftop for solar drift-scan observations at four east–west baselines. This section describes the antenna setup, front-end configuration, data acquisition pipeline, and the observing sessions used in the analysis presented in Chapter 5.

4.3.1 Observation Setup

All observations were carried out in drift-scan mode at 1.420 GHz. The antennas were fixed in azimuth and elevation, and the Sun drifted through the beam as the Earth rotated.

Antenna Pointing and Site Evolution

Both antennas were pointed due south (180° azimuth). The elevation for each session was set from the expected solar transit altitude, computed using

$$\sin(A_{\rm el}) = \sin\delta_\odot \, \sin\phi + \cos\delta_\odot \, \cos\phi \, \cos H, \tag{4.3}$$

which reduces at meridian transit ($H = 0$) to

$$A_{\rm el} = 90^\circ - |\phi - \delta_\odot|, \tag{4.4}$$

where $A_{\rm el}$ is the elevation, $\phi = 22.7^\circ$ N is the latitude of IIT Indore and $\delta_\odot$ is the solar declination on the date of observation. The antennas were manually set to the elevation given by Eq. (4.4) before each session.

The observation campaign began in March 2026 from a covered balcony. At this time $\delta_\odot \approx 0^\circ$ (the vernal equinox falls on 20 March 2026), so the predicted elevation at transit was $A_{\rm el} = 90^\circ - 22.7^\circ \approx 67.3^\circ$. An overhead structure on the balcony limited the maximum reachable elevation to $\sim 60^\circ$, placing the Sun above the beam centre at transit. This caused partial obstruction, reduced fringe clarity, and increased system noise in the early data.

To overcome this constraint, the system was relocated to an open rooftop for all subsequent sessions. Observations were conducted in early April 2026, when $\delta_\odot \approx +3^\circ$ to $+7^\circ$, giving a predicted transit altitude of approximately 75° from Eq. (4.4). The rooftop allowed the antennas to be set to this elevation without obstruction, keeping the Sun within the main beam at transit. This produced a clear improvement in fringe stability and signal-to-noise ratio, and all four sessions retained for analysis were conducted in this configuration. For each session, the actual value of $\delta_\odot$ on the observation date was used in both the antenna pointing calculation and the fringe period prediction (Eq. 2.17), rather than a single fixed approximation.

The uncertainty in baseline measurement contributes only a small error to the predicted fringe period. Since

$$T_{\rm pred} \propto \frac{1}{B},$$

the fractional uncertainty is

$$\frac{\delta T_{\rm pred}}{T_{\rm pred}} = \frac{\delta B}{B}.$$

For a baseline uncertainty of $\delta B = \pm 5\,\text{mm}$, this corresponds to a fractional uncertainty of approximately 0.5% for the 1.00 m baseline and 0.2% for the 2.40 m baseline. These values are significantly smaller than the observed $\sim$ 16% discrepancy between measured and predicted fringe periods, indicating that baseline measurement uncertainty is not the dominant source of error.



(a) Balcony setup with limited elevation ($\sim 60°$).

(b) Rooftop setup with improved sky access.

Figure 4.17: Comparison of observing configurations. Moving to the rooftop improved data quality and reduced obstruction effects.

4.3.2 Front-End Signal Chain for Outdoor Observations

All observations used the final front-end configuration described in Chapter 3. Identical RF chains were used for both antennas to maintain symmetry between the two signal paths.

Matching the front-end chains is important because any mismatch can introduce artificial amplitude or phase differences. Both chains were characterised with a Vector Network Analyser before deployment.

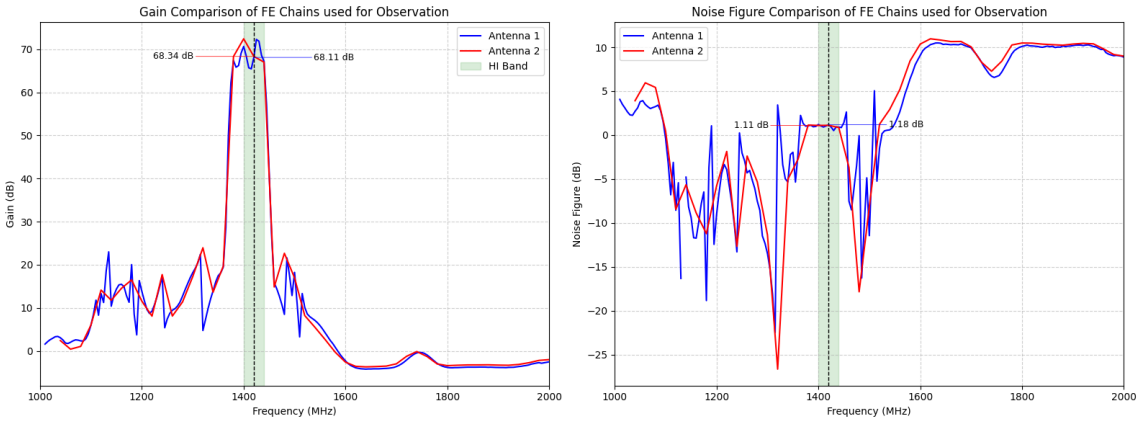

Figure 4.18: Measured gain (left) and noise figure (right) for the two front-end chains. The shaded region marks the H I band around 1.42 GHz. Both chains show closely matched responses within the observation band.

The selected configuration was designed to provide a total gain of ~ 80 dB. During the VNA characterisation measurements shown in Figure 4.18, an external attenuator was inserted to prevent receiver saturation. In addition, the measurements included the actual cables and connectors used during the observations. As a result, the measured end-to-end gain within the H I band was approximately ~ 68 dB. Within the H I band, both chains show closely matched gain (~ 68 dB) and noise figure (~ 1.1–1.2 dB), helping ensure that measured visibilities are dominated by the sky signal rather than by instrumental imbalance.

4.3.3 Digital Back-End: FX Correlator

The digital backend uses an FX correlator for real-time estimation of complex visibility. The GNU Radio flowgraph used for all observations is shown in Figure 4.19.

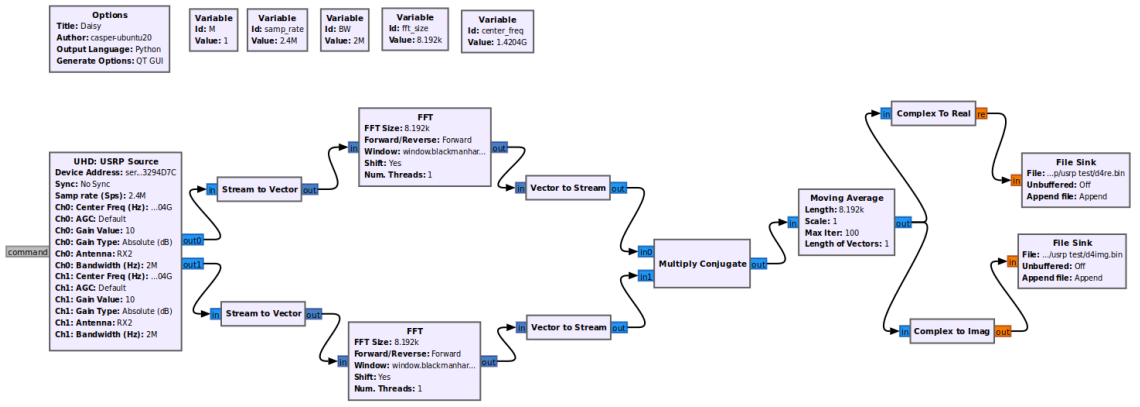

Figure 4.19: GNU Radio FX correlator pipeline. The two input channels are Fourier transformed and cross-correlated to obtain complex visibility.

Both channels were sampled at 2 Msps and centred at 1420 MHz. The samples were grouped

into blocks, Fourier transformed, and multiplied by the complex conjugate of the second channel to form the cross-power spectrum. A moving average was applied to improve the signal-to-noise ratio, and the resulting complex visibility was recorded for offline analysis.

Observing Schedule and Session Duration

Each session was centred on solar transit, typically from 11:00 to 14:00 IST, giving a total duration of 3 hours. This provided full coverage of the solar passage through the beam, along with baseline measurements before and after transit.

Baseline Configuration

The projected antenna separation defines the interferometric baseline B, which was varied between sessions to examine its effect on fringe behaviour. The baselines used in this work were 1.00 m, 1.88 m, 2.00 m, and 2.40 m. At 1.420 GHz ($\lambda \approx 0.211$ m), these correspond to 4.74 λ, 8.9 λ, 9.5 λ, and 11.4 λ, respectively. The dependence of fringe properties on baseline length is discussed in Chapter 5.

Chapter 5

Results

5.1 RTL-SDR Backend: Phase Coherence Test

5.1.1 Cross-Correlation Result

The RTL-SDR bench test described in Section 4.1 was the first cross-correlation measurement performed in this project. The two daisy-chained dongles received the same split tone, and the cross-power spectrum was computed continuously over the full observation window.

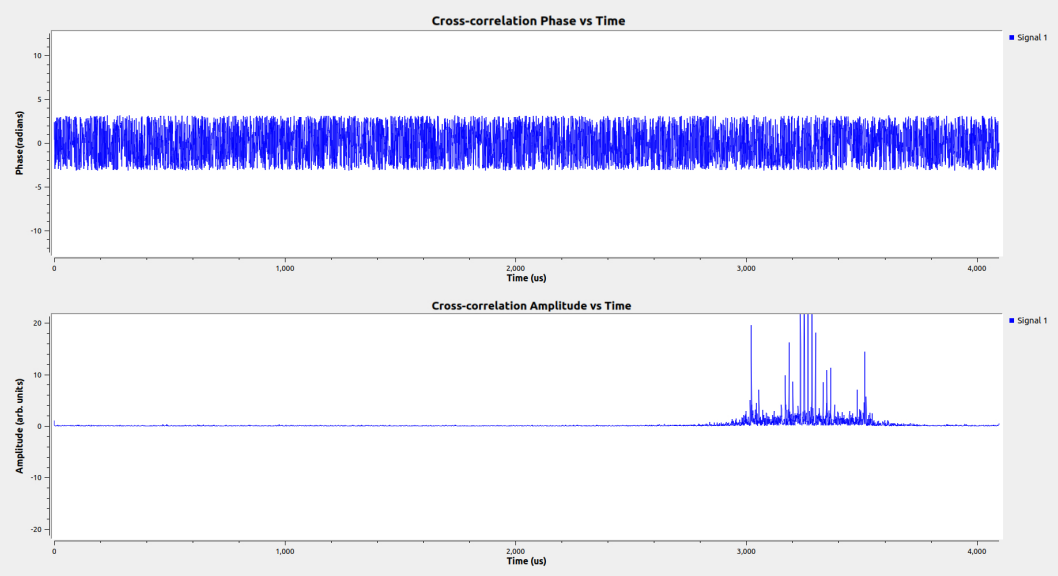

Figure 5.1: Cross-correlation result for two daisy-chained RTL-SDR dongles receiving the same split tone. The phase (upper) spans $[-\pi, +\pi]$ with no stable structure, while the amplitude (lower) remains near zero except for a brief startup transient around $t \approx 2800$–$3400\,\mu s$. No stable cross-correlation phase is established, showing that the daisy-chain provides a common reference frequency but not a shared phase.

The phase panel shows no preferred value and no tendency to stabilise at a fixed offset. This behaviour is characteristic of two independent oscillators whose relative phase drifts continuously. As a result, the cross-product $V_1 \cdot V_2^*$ accumulates phase at a rate set by the frequency difference

between the two PLLs, which is neither zero nor constant. The amplitude burst visible near $t \approx 3000$ µs is a brief transient at startup and disappears quickly, leaving no sustained correlated signal.

This result established the requirements for every subsequent backend. A useful interferometer must produce a stable, non-zero amplitude and a flat or slowly varying phase in the cross-correlation output. The RTL-SDR daisy-chain satisfies neither condition.

5.1.2 Limitations of Daisy-Chain

The failure is instructive and worth stating precisely, because it directly motivated the design choices made for the USRP B210 backend. The daisy-chain routes a common 28.8 MHz reference frequency from the master dongle to the slave, so both PLLs are frequency-locked. However, each PLL settles to a random initial phase at power-on, and the phase noise accumulated by each PLL over time is statistically independent. The result is an inter-device phase offset that is random at startup and continues to drift throughout the observation.

The contrast with the USRP B210 is direct. In the B210, both receive channels are processed by the same AD9361 chip, driven by the same internal oscillator and the same PLL. There is one clock, not two. The phase relationship between the two channels is therefore fixed from power-on and does not drift, as confirmed experimentally in Section 5.2.1.

5.2 USRP B210 Backend: Phase Synchronisation Results

This section presents the phase synchronisation results obtained with the USRP B210 backend. The results are discussed in three stages: first, the internal phase coherence of a single device; second, synchronisation between two independent units; and finally, simultaneous operation across all six baselines of the four-element system. All measurements were carried out using GNU Radio flowgraphs with the synchronisation modifications to the generated Python code described in Chapter 4.

5.2.1 Stage 1: Single USRP B210 — Internal Phase Coherence

The first experiment examined the phase relationship between the two receive channels of a single USRP B210. A continuous-wave tone at 1420.405 MHz was generated and split using a two-way power divider, providing identical signals to both RX2 ports. The resulting streams were processed using the FX correlator described in Chapter 4, and the cross-product $V_0 \cdot V_1^*$ was computed.

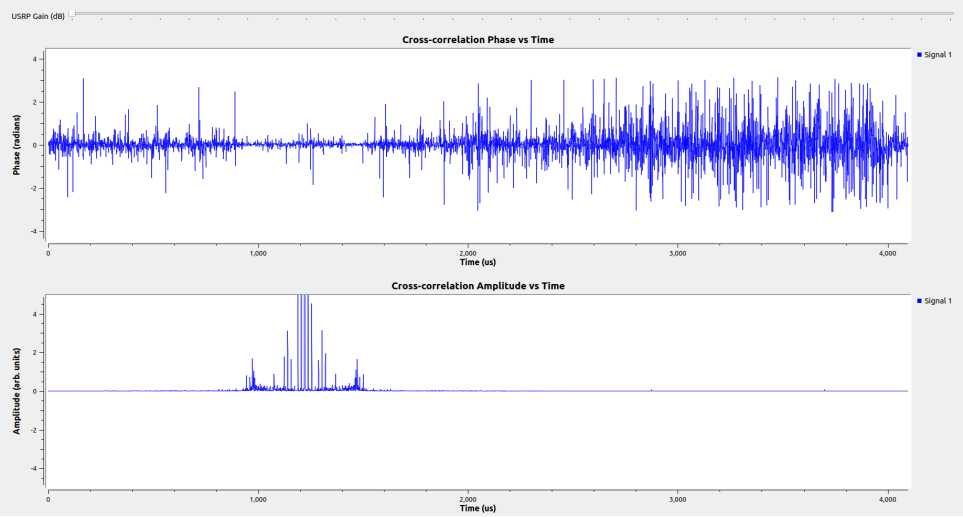

Figure 5.2: Single USRP B210 internal coherence test. The upper panel shows the instantaneous cross-correlation phase, while the lower panel shows the corresponding amplitude. A strong correlated signal appears as a distinct peak in amplitude. The phase fluctuates because of thermal noise and finite signal-to-noise ratio, but remains statistically stable over time.

The cross-correlation amplitude shows a clear peak at the injected tone, indicating that both channels receive the same signal. The instantaneous phase is not constant and exhibits rapid fluctuations, which is expected because the phase estimate is sensitive to noise at the sample level. When the signal amplitude approaches the noise floor, the phase becomes poorly defined and appears random.

No systematic drift is observed over time. This shows that the relative phase between the two channels remains stable, since both share a common local oscillator and RF chain within the AD9361. When averaged over time, the phase converges to a constant value, demonstrating internal phase coherence. These results show that a single USRP B210 can provide a stable two-channel interferometric measurement and establish a baseline for extension to multiple devices.

5.2.2 Stage 2: Two-USRP Inter-Device Synchronisation

Once single-device coherence had been confirmed, the next step was to establish phase-coherent operation between two independent USRP B210 units. Unlike the single-device case, each unit has its own local oscillator, so a stable phase relationship cannot be maintained without external synchronisation.

Baseline: No External Synchronisation

To demonstrate this limitation, both USRPs were driven by a common 1420.405 MHz tone while operating on their internal clocks, with no external reference or PPS alignment.

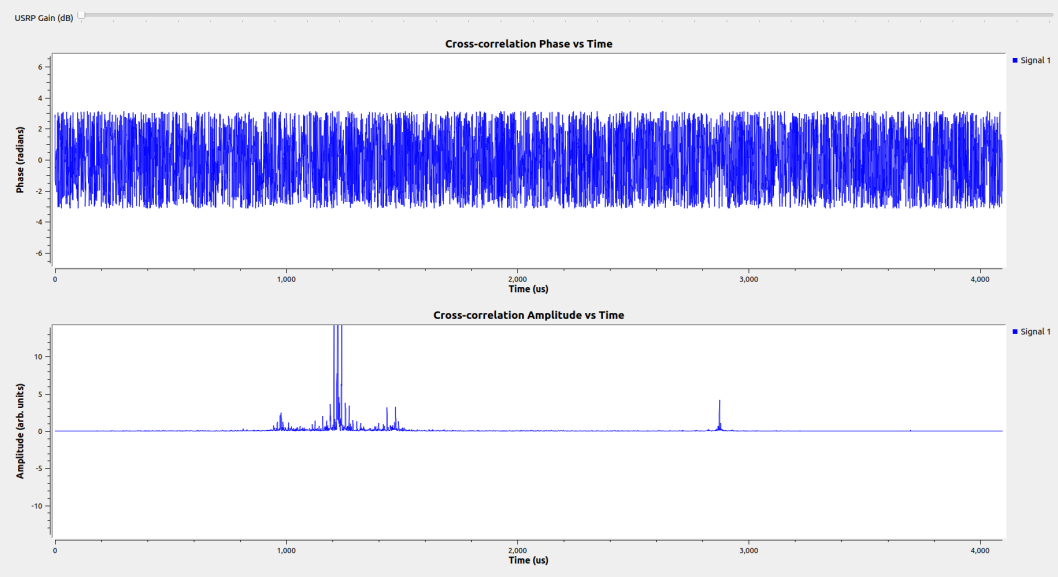

Figure 5.3: Two-USRP cross-spectrum without external synchronisation. The cross-correlation phase (upper panel) is distributed across $[-\pi, +\pi]$ with no stable structure, indicating complete phase incoherence. The amplitude (lower panel) shows only weak and inconsistent correlation, confirming that the two devices are not phase-locked.

The absence of a stable phase structure shows that independent oscillators introduce a time-varying phase, making long-term interferometric measurements unreliable.

Time-Domain Result With Synchronisation

External synchronisation was implemented using a shared 10 MHz reference and PPS signal, together with PPS-aligned timed-start control in software (Chapter 4). Under these conditions, a stable phase relationship was recovered.

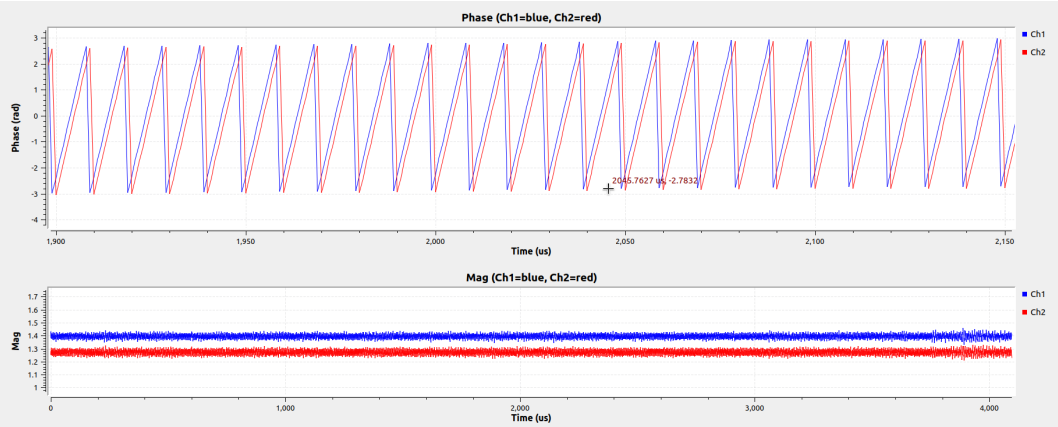

Figure 5.4: Two-USRP cross-correlation with full synchronisation. The phase (upper panel) remains nearly constant at ~ -2.8 rad with only small fluctuations, indicating successful phase locking. The magnitude (lower panel) is stable and non-zero, confirming continuous correlated signal detection.

Frequency-Domain Cross-Spectrum

The synchronised system was also verified in the frequency domain using an FX correlator.

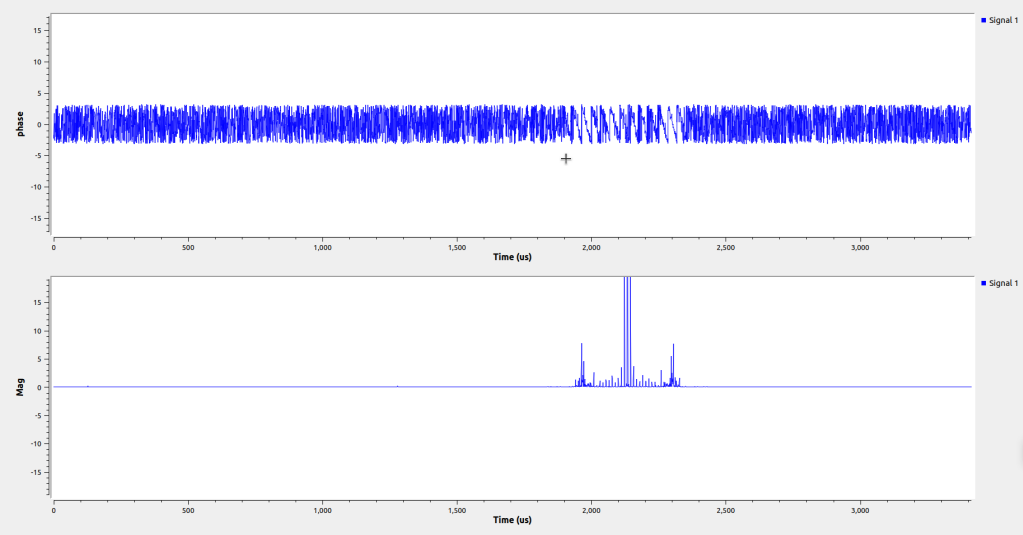

Figure 5.5: FFT cross-spectrum with synchronisation enabled. The phase (upper panel) remains bounded with no systematic drift, while the amplitude (lower panel) shows a clear correlated peak at the injected tone frequency. The noise bins remain uncorrelated, confirming that the observed structure arises from the signal rather than from system artefacts.

The presence of a distinct spectral peak with stable phase confirms that coherence is maintained across the full bandwidth.

5.2.3 Physical Interpretation of the Phase Behaviour

Fixed LO Phase Offset

All synchronised measurements show a constant, non-zero phase offset. This arises at power-on when the two AD9361 PLLs lock to the shared 10 MHz reference. Although the frequencies are identical, the initial phase is random, producing a fixed offset for each session. This offset remains stable over time and can be removed through standard calibration. In our measurements (Figure 5.4), this offset was approximately 2.8 rad, which remained stable throughout the session.

Phase Behaviour in Time Domain

The individual channel phases show a sawtooth pattern because of a small frequency offset between the received tone and the digital downconversion frequency. The phase therefore increases linearly and wraps at $\pm\pi$. The key observable is the phase difference. When both devices are synchronised, their phase evolution is identical and the phase difference remains constant. Without synchronisation, different phase rates lead to continuous drift, which prevents interferometric operation.

5.2.4 Stage 3: Four-Element Six-Baseline Validation

Inter-device synchronisation was established on a single cross-device baseline before extending the system to all four channels by enabling both receive channels on each USRP. The correlator then computed all six cross-products simultaneously using the same synchronisation scheme.

The four channels are labelled A1 and A2 (USRP 1, channels 0 and 1) and B1 and B2 (USRP 2, channels 0 and 1). The resulting six baselines are listed in Table 4.2. Baselines A1–A2 and B1–B2 are internal and therefore inherently phase-coherent, whereas the remaining four are cross-device baselines that depend on external synchronisation.

Without External Reference

Before synchronisation was applied, the four-element correlator was operated without an external 10 MHz reference or PPS signal to establish the incoherent baseline case.

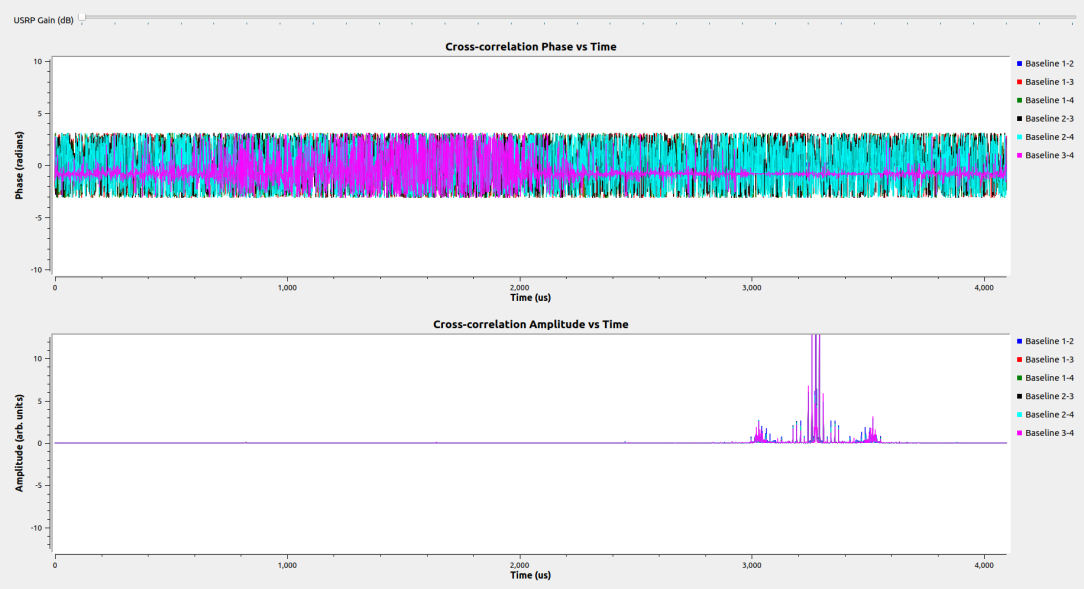

Figure 5.6: Four-element, six-baseline cross-correlation without external synchronisation. The phase (top) is erratic across all baselines, with no stable structure, while the amplitude (bottom) remains coherent only for the internal Baseline 3–4 within USRP 2. All cross-device baselines are incoherent, confirming the need for external synchronisation.

All Six Baselines With Synchronisation: Overview

Figure 5.7 shows the cross-correlation phase and amplitude for all six baselines under full synchronisation.

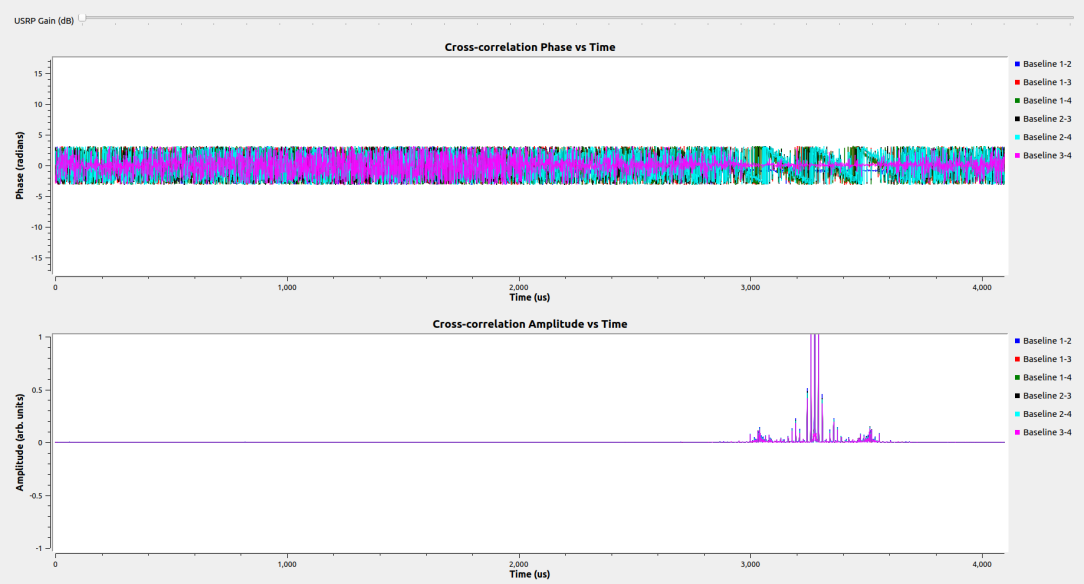

Figure 5.7: Overview of all six baselines with synchronisation active. Phase (upper) and amplitude (lower) versus time. Colour coding: 1–2 (blue), 1–3 (red), 1–4 (green), 2–3 (black), 2–4 (cyan), 3–4 (magenta). All phase traces remain bounded and stable, and correlated amplitude spikes appear simultaneously across all baselines near $t \approx 3000$ µs.

In contrast to Figure 5.6, the bounded and structured phase behaviour across all six baselines, including the cross-device baselines, arises directly from the external synchronisation procedure.

Individual Baseline Results

Each baseline was also examined individually to assess the quality of the phase behaviour in isolation.

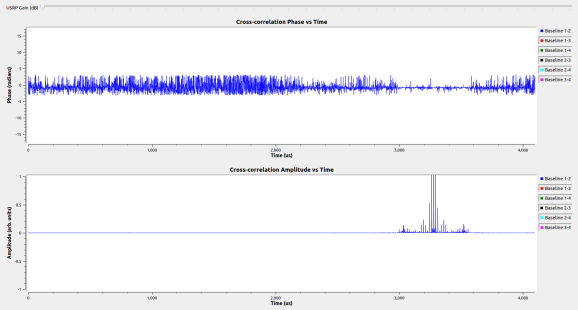

(a) Baseline 1–2 (A1×A2, internal to USRP 1).

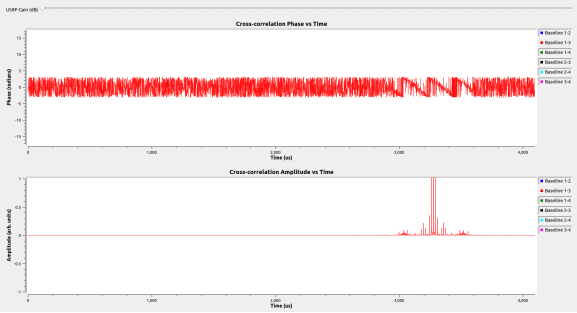

(b) Baseline 1–3 (A1×B1, cross-device).

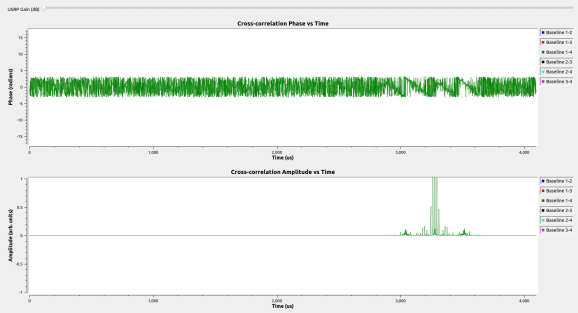

(c) Baseline 1–4 (A1×B2, cross-device).

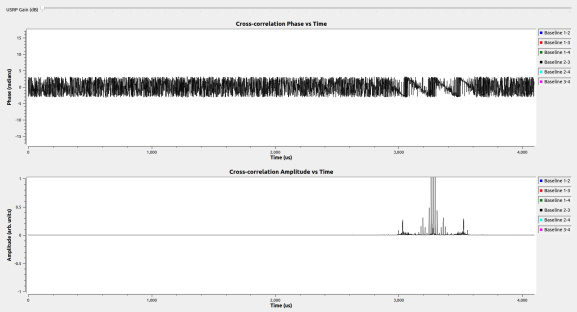

(d) Baseline 2–3 (A2×B1, cross-device).

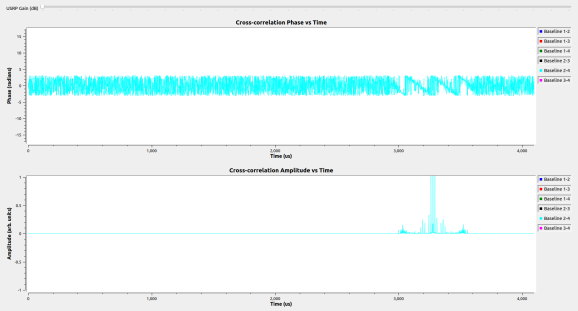

(e) Baseline 2–4 (A2×B2, cross-device).

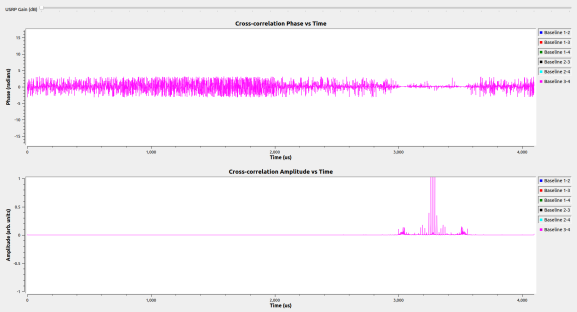

(f) Baseline 3–4 (B1×B2, internal to USRP 2).

Figure 5.8: Individual cross-correlation phase (upper panel) and amplitude (lower panel) for all six baselines. The two internal baselines (1–2 and 3–4) show the cleanest phase behaviour, as both channel pairs share a single AD9361 chip. All four cross-device baselines (1–3, 1–4, 2–3, 2–4) show bounded phase with no sweeping and correlated amplitude spikes appearing in the same ~2800–3400 µs window, confirming phase-coherent operation across every inter-device channel combination.

Interpretation

In all six panels the cross-correlation phase is contained within $[-\pi, +\pi]$ with no secular drift, in direct contrast to the incoherent case. The amplitude panels show near-zero response everywhere except for a cluster of correlated spikes appearing in approximately the same 2800–3400 µs window across all baselines simultaneously. Since all six baselines receive the same injected tone, this simultaneous correlated response confirms that the four-element system is behaving as

a single coherent instrument.

The cross-device baselines show slightly greater phase scatter than the internal baselines, consistent with residual timing jitter introduced by the Python GIL during the threaded `set_time_next_pps()` calls (Section 4.2.4). This scatter does not indicate drift and does not prevent the detection of correlated signals.

5.2.5 USRP Synchronisation Summary

The results point to three main conclusions. A single USRP B210 provides phase-coherent operation between its two receive channels without additional hardware. Two independent B210 units can also be made phase-coherent by using a shared 10 MHz reference and PPS signal together with a PPS-aligned timed start in software. This method extends directly to all four channels and six baselines, as confirmed by the four-element tests. A session-dependent LO phase offset remains after each power cycle, but it can be removed through standard calibration and does not limit the usefulness of the system.

5.2.6 Representative Solar Interferometric Fringe Observations

Four observing sessions were conducted on the IIT Indore rooftop during April 2026, each lasting three hours and centred on solar transit. The antennas were fixed at azimuth 180° (due south) and elevation $\approx 70°$, with the Sun drifting through the beam as the Earth rotated. East–west baselines of 1.00, 1.88, 2.00, and 2.40 m were used in successive sessions. The results are presented in order of increasing baseline length.

As noted in Section 2.5, the fringe period is not strictly constant during a drift scan because it depends on $\cos H$, which decreases as the source moves away from the meridian. Since each session spans ± 1.5 hours ($H_{\max} = 22.5°$) around transit, the predicted fringe period is best compared with the session-averaged measured period rather than the instantaneous transit value. The time-averaged $\langle \cos H \rangle$ over the observing window is:

$$\langle \cos H \rangle = \frac{\sin(H_{\max})}{H_{\max}} = \frac{\sin(22.5°)}{22.5° \times \pi/180} = \frac{0.3827}{0.3927} \approx 0.974, \tag{5.1}$$

and the session-averaged predicted fringe period is obtained by substituting this into Eq. (2.18):

$$\langle T_{\text{fringe}} \rangle = \frac{(\lambda/B) \cdot (180/\pi)}{0.25 \cos \delta_{\odot} \langle \cos H \rangle}, \tag{5.2}$$

where $\cos \delta_{\odot} = 0.998$ corresponds to the solar declination $\delta_{\odot} \approx 3°$–$7°$ during the April 2026 sessions, and $\lambda = 0.211$ m at 1.420 GHz. This averaged prediction is used throughout the subsections below.

5.2.7 Baseline $B = 1.00$ m ($u = 4.74\,\lambda$)

At the shortest baseline, the fringe spacing on the sky is the largest ($\theta_f = \lambda/B = 12.1°$), and the interferometer is least sensitive to fine angular structure in the source. Applying Eq. (5.2):

$$\langle T_{\text{fringe}} \rangle = \frac{(0.211/1.00) \cdot (180/\pi)}{0.25 \times 0.998 \times 0.974} \approx 49.8 \text{ min}. \tag{5.3}$$

Over the 3-hour session the expected fringe count is $180/49.8 \approx 3.6$, and three peaks were observed, in good agreement. The measured period, $T_{\text{meas}} = 44.4 \pm 2.8$ min, is shorter than the prediction by 10.9%. As discussed in Section 5.3, the measured period tends to be biased toward the transit value because fringe peaks are most clearly resolved near transit where the signal-to-noise ratio is highest; this systematically shortens the measured average relative to the true session-averaged prediction. The median visibility magnitude was 29.9 dB and the phase standard deviation was 0.132 rad, both indicating a stable and well-detected fringe pattern at this baseline.

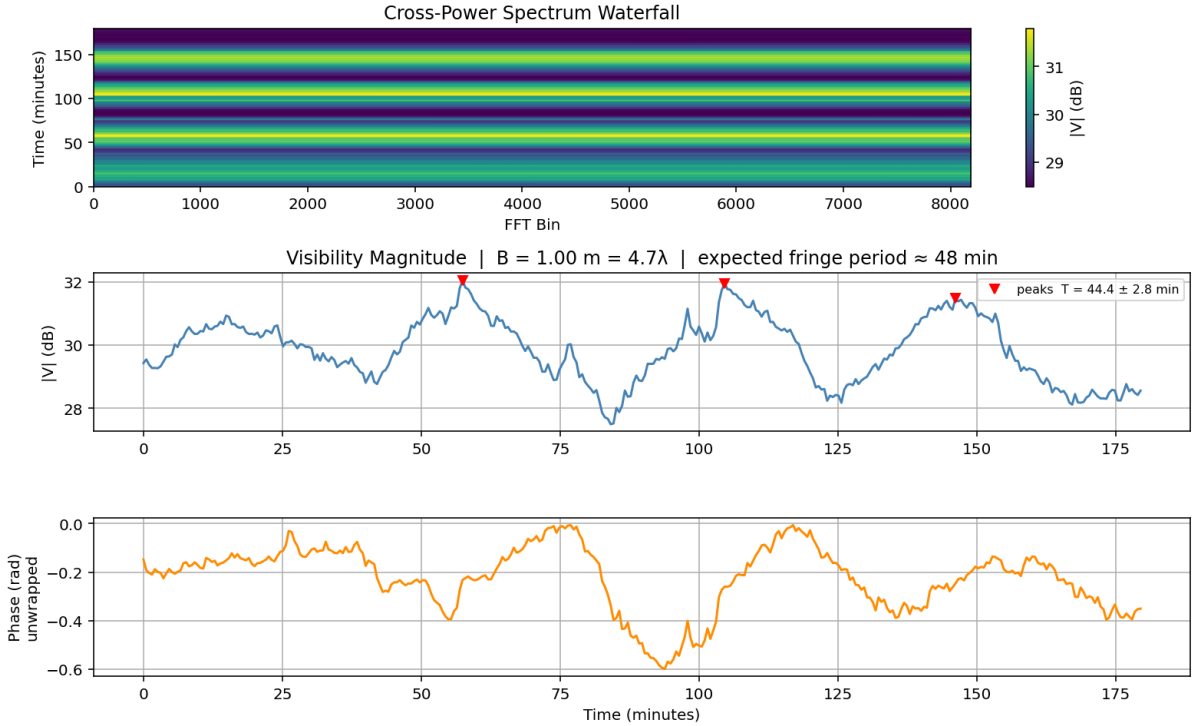

Figure 5.9: Solar drift-scan observation at baseline $B = 1.00$ m ($u = 4.74\,\lambda$). The upper panel shows the cross-power spectrum waterfall over the 3-hour session. The middle panel shows the visibility magnitude as a function of time, with three fringe peaks clearly identified (marked by triangles); the measured period is 44.4 ± 2.8 min. The lower panel shows the corresponding unwrapped visibility phase.

5.2.8 Baseline $B = 1.88$ m ($u = 8.9\,\lambda$)

Increasing the baseline to 1.88 m almost doubles the fringe rate, reducing the fringe spacing to $\theta_f = \lambda/B = 6.44°$. Applying Eq. (5.2):

$$\langle T_{\text{fringe}} \rangle = \frac{(0.211/1.88) \cdot (180/\pi)}{0.25 \times 0.998 \times 0.974} \approx 26.5 \text{ min.} \tag{5.4}$$

The expected fringe count over 3 hours is $180/26.5 \approx 6.8$, and six peaks were observed, which is consistent with this prediction. The measured period, $T_{\text{meas}} = 25.6 \pm 1.5$ min, agrees with the predicted value to within 3.4%, making this the closest match among all four sessions and well within the measurement uncertainty. This level of agreement suggests that the atmospheric and instrumental conditions were particularly stable during the observation. The median visibility magnitude was 25.0 dB and the phase standard deviation was 0.330 rad. Although the phase scatter is somewhat higher than in the 1 m session, it is more likely to reflect short-timescale atmospheric fluctuations than any loss of instrumental coherence.

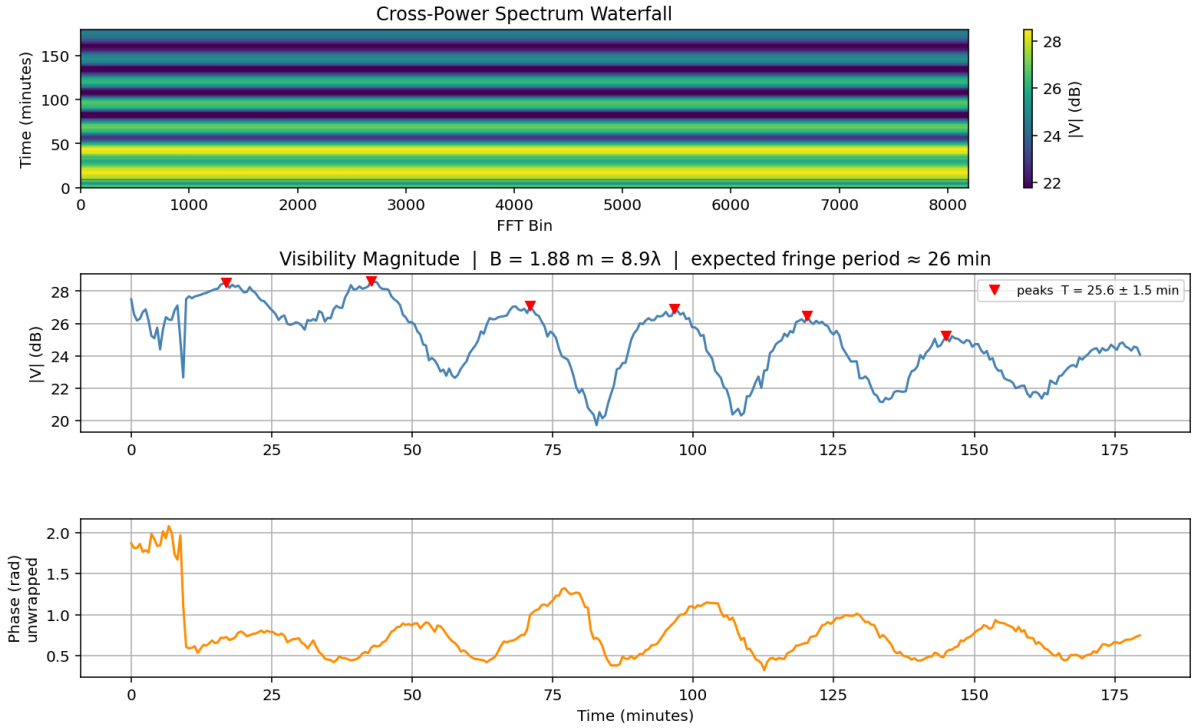

Figure 5.10: Solar drift-scan observation at baseline $B = 1.88$ m ($u = 8.9\,\lambda$). Six fringe peaks are identified over the 3-hour interval, giving a measured period of 25.6 ± 1.5 min, in agreement with the session-averaged prediction of 26.5 min to within 3.4%.

5.2.9 Baseline $B = 2.00$ m ($u = 9.5\,\lambda$)

At 2.00 m separation the fringe spacing narrows to $\theta_f = \lambda/B = 6.05°$. Applying Eq. (5.2):

$$\langle T_{\text{fringe}} \rangle = \frac{(0.211/2.00) \cdot (180/\pi)}{0.25 \times 0.998 \times 0.974} \approx 24.9 \text{ min.} \tag{5.5}$$

The expected fringe count for this session is $180/24.9 \approx 7.2$, while six peaks were observed. The measured period, $T_{\text{meas}} = 23.2 \pm 1.8$ min, agrees with the predicted value to within 7.3%. This was the most stable of the four sessions, with the highest median visibility magnitude of

31.5 dB and the lowest phase standard deviation of 0.087 rad, indicating strong instrumental coherence and very limited atmospheric disturbance throughout the three-hour observation. The shortfall of one fringe relative to the expected count is most likely explained by a peak falling below the detection threshold near the edges of the observing window, where the Sun is furthest from the beam centre and the signal-to-noise ratio is lowest.

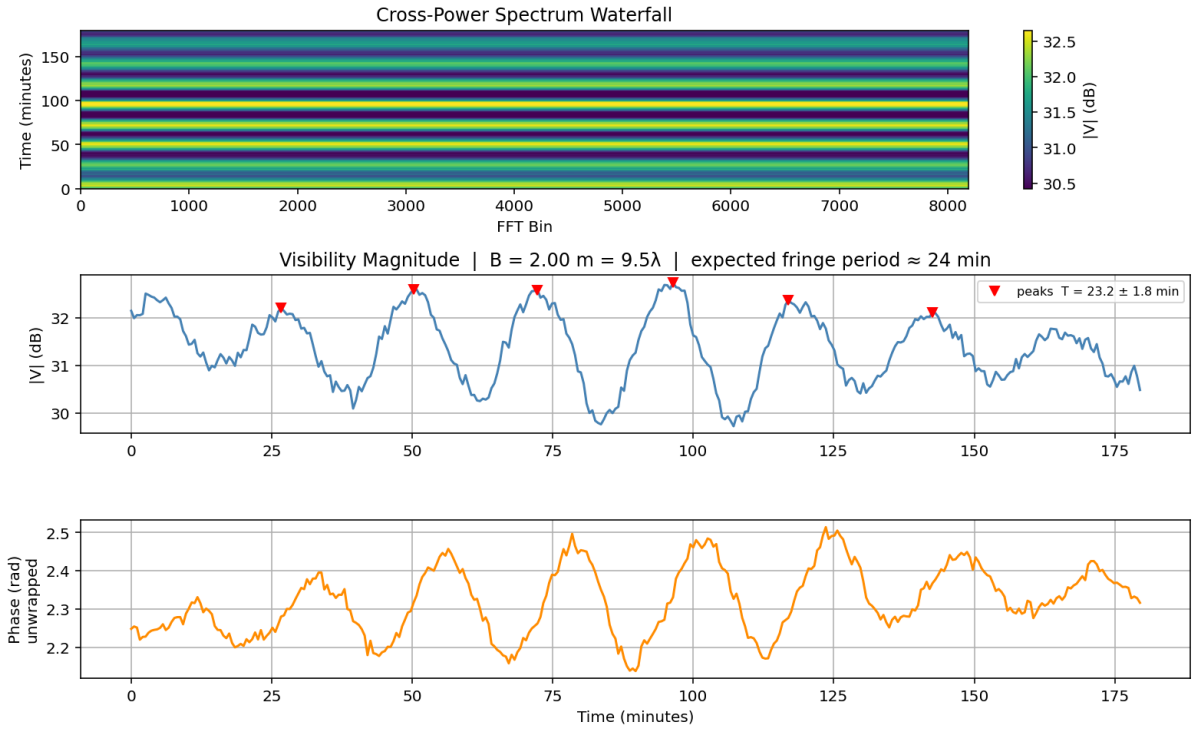

Figure 5.11: Solar drift-scan observation at baseline $B = 2.00$ m ($u = 9.5\,\lambda$). Six fringe peaks are identified over the 3-hour session, giving a measured period of 23.2 ± 1.8 min compared with the session-averaged prediction of 24.9 min (agreement within 7.3%). This session gave the highest median visibility magnitude (31.5 dB) and the lowest phase scatter (0.087 rad) of all four sessions.

5.2.10 Baseline $B = 2.40$ m ($u = 11.4\,\lambda$)

The longest baseline used in this study, at 2.40 m, produces the finest fringe spacing of $\theta_f = \lambda/B = 5.04°$ and the highest fringe rate among the four sessions. Applying Eq. (5.2):

$$\langle T_{\text{fringe}} \rangle = \frac{(0.211/2.40) \cdot (180/\pi)}{0.25 \times 0.998 \times 0.974} \approx 20.8 \text{ min}. \tag{5.6}$$

The expected fringe count is $180/20.8 \approx 8.7$, and the nine observed peaks are consistent with this prediction. The measured period, $T_{\text{meas}} = 17.5 \pm 2.0$min, is shorter than the predicted value by $\sim 16\%$, representing the largest discrepancy among the four sessions. Even so, the detection of nine peaks, the highest count among all sessions, clearly supports the expected increase in fringe rate with baseline length. The median visibility magnitude was 24.1 dB and the phase standard deviation was 0.330 rad, both comparable to the 1.88 m session and consistent with

stable system performance.

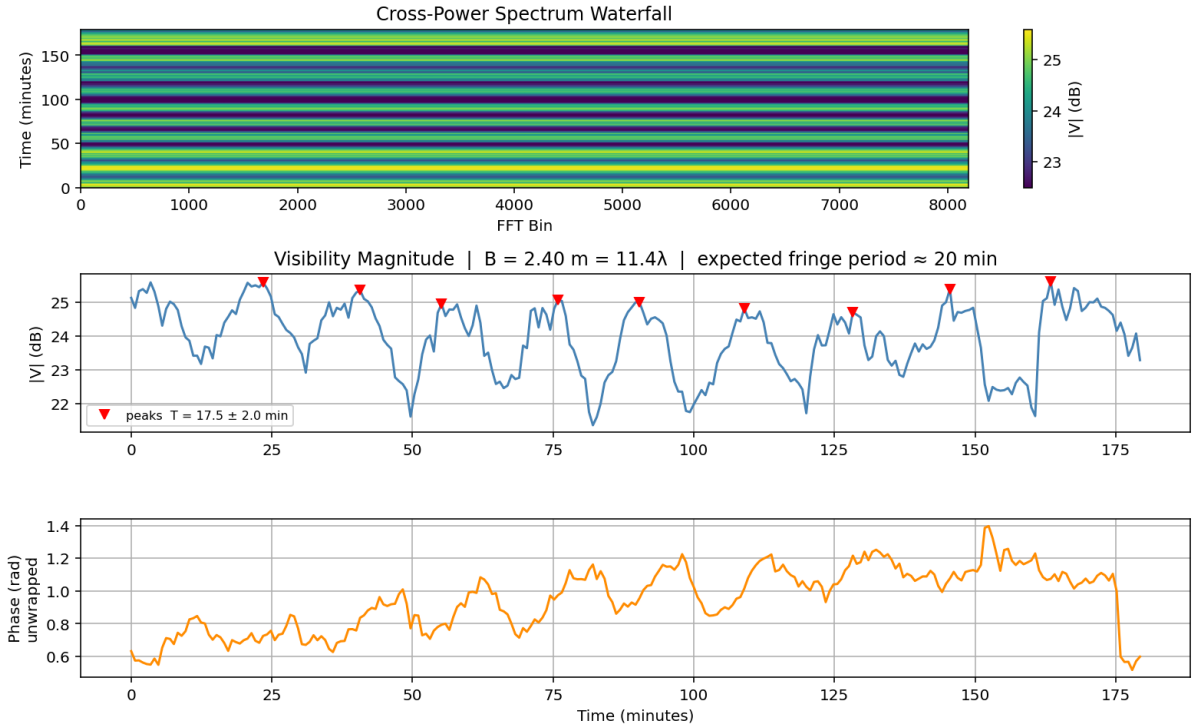

Figure 5.12: Solar drift-scan observation at $B = 2.40$ m ($u = 11.4\ \lambda$). Nine fringe peaks are detected over 3 hours, consistent with the expected count of 8.7. The measured period, 17.5 ± 2.0 min, is $\sim 16\%$ shorter than the predicted 20.8 min, likely due to SNR-related peak selection bias and increased baseline uncertainty at longer separations.

5.3 Dependence of Fringe Properties on Baseline

The four observing sessions provide a controlled framework for examining how interferometric fringe properties vary with baseline length. Since the only parameter changed between sessions was the antenna separation, while the observing frequency, site, and signal chain were kept fixed, the results allow the effect of baseline on the measured visibility to be isolated more clearly. This section focuses on three independently measurable quantities: fringe period, fringe count, and visibility amplitude, and compares the observed behaviour with the predictions of the interferometric model developed in Chapter 2.

Fringe Period and Its Baseline Dependence

The session-averaged fringe period is predicted by Eq. (5.2), which gives an inverse relationship between $\langle T_{\text{fringe}} \rangle$ and B. Table 5.1 summarises the predicted and measured periods for all four sessions. The measured periods decrease monotonically from 44.4 ± 2.8 min at $B = 1.00$ m to 17.5 ± 2.0 min at $B = 2.40$ m, consistent with the predicted inverse scaling.

The residual discrepancies range from 3.4% at $B = 1.88$ m to $\sim 16\%$ at $B = 2.40$ m. A systematic contribution to these discrepancies arises from the SNR-driven peak selection bias: fringe peaks are most reliably identified near transit, where the Sun is closest to the beam centre and the signal-to-noise ratio is highest. As a result, the measured average period is weighted toward the transit value, which is shorter than the true session average by approximately $1/\langle \cos H \rangle - 1 \approx 2.7\%$. This systematic offset is present in all sessions and accounts for part of the discrepancy at every baseline. The remaining scatter is consistent with peak identification uncertainty in the presence of noise and the ±5 mm uncertainty in baseline measurement.

Table 5.1: Summary of fringe parameters for the four representative observation sessions. Each session had a duration of 3 hours at azimuth 180° and elevation ≈ 70°. Predicted periods use the session-averaged hour angle correction $\langle \cos H \rangle = 0.974$ (Eq. 5.1). Agreement is defined as $|T_{\text{meas}} - \langle T_{\text{pred}} \rangle| / \langle T_{\text{pred}} \rangle \times 100\%$.

B (m)	u (λ)	$\langle T_{\text{pred}} \rangle$ (min)	T_{meas} (min)	Agree. (%)	Peaks (detected)	Med. vis. (dB)	σ_ϕ (rad)
1.00	4.74	49.8	44.4 ± 2.8	10.9	3	29.9	0.132
1.88	8.90	26.5	25.6 ± 1.5	3.4	6	25.0	0.330
2.00	9.50	24.9	23.2 ± 1.8	7.3	6	31.5	0.087
2.40	11.40	20.8	17.5 ± 2.0	15.9	9	24.1	0.330

Fringe Rate and Peak Count

An independent check of the fringe period measurements can be made by counting the number of fringe peaks detected during the fixed 3-hour observing window. If the fringe rate scales as $1/B$, then the peak count should increase linearly with baseline length. The observed counts of 3, 6, 6, and 9 for baselines of 1.00, 1.88, 2.00, and 2.40m, respectively, increase monotonically with B, which is consistent with this expectation. The counts predicted by Eq.(5.2) are 3.6, 6.8, 7.2, and 8.7, and each differs from the observed value by no more than one peak. This agreement is particularly encouraging because peak counting is completely independent of the period measurement. It does not require any assumption about the detailed shape of the fringe waveform, but only that individual maxima can be identified above the noise floor. The consistency between these two methods therefore provides added confidence in the robustness of the underlying fringe detection.

Visibility Amplitude and Source Structure

The amplitude of the measured visibility contains information about the angular structure of the source on the scale probed by the baseline. For a uniform circular disc of angular diameter θ_s, the visibility amplitude follows a first-order Bessel profile and becomes zero at the null baseline

$B_{\text{null}} = 1.22\lambda/\theta_s$. For the Sun at 1.42 GHz, with a disc diameter of $\theta_s \approx 0.5°$, this gives:

$$B_{\text{null}} = \frac{1.22 \times 0.211}{0.5° \times \pi/180} \approx 29.5 \text{ m}. \tag{5.7}$$

All four baselines used here (1.00 – 2.40 m) are far below this value, which means that the Sun remains completely unresolved at these spacings. The visibility amplitude is therefore expected to be effectively independent of baseline. The measured median visibility magnitudes of 29.9, 25.0, 31.5, and 24.1 dB for baselines of 1.00, 1.88, 2.00, and 2.40 m, respectively, show no systematic trend with B, in line with this expectation. The scatter of ~ 7 dB across sessions is more likely to reflect day-to-day gain variations in the front-end chains, differences in the level of radio-frequency interference, and the absence of bandpass calibration, rather than any real change in solar structure. The session with $B = 2.00$ m produced both the highest visibility magnitude and the lowest phase scatter, suggesting that it was recorded under the most stable instrumental and atmospheric conditions.

Phase Stability and Coherence

The phase standard deviation σ_ϕ listed in Table 5.1 measures how steadily the visibility phase evolves over the session. A small σ_ϕ indicates that the fringe phase changes smoothly and predictably, as expected for a point source observed under stable conditions. The lowest value, $\sigma_\phi = 0.087$ rad at $B = 2.00$ m, is consistent with the high visibility magnitude and confirms that the signal chain was operating coherently throughout that session. The sessions at $B = 1.88$ m and $B = 2.40$ m both show $\sigma_\phi = 0.330$ rad, roughly four times larger, suggesting stronger short-timescale atmospheric phase fluctuations on those days rather than any instrumental problem, since the visibility amplitudes remain stable in both cases. At $B = 1.00$ m the intermediate value of 0.132 rad reflects the slower phase evolution at the longest fringe period, which reduces the instantaneous phase noise even under moderately variable conditions.

Across all sessions, no secular phase drift was observed. The phase remained bounded and varied smoothly on timescales consistent with the fringe period, confirming that the USRP B210 backend maintained phase coherence throughout each three-hour observation without the need for mid-session recalibration.

Chapter 6

Summary and Future Scope

6.1 Summary of the Project

This thesis presents the design, construction, and testing of a two-element radio interferometer operating at 1.420 GHz, the neutral hydrogen line. The work was motivated by the limited availability of simple, well-documented teaching interferometers, which often leaves students with only a theoretical understanding of the technique.

6.1.1 Instrument Design and System Architecture

The instrument consists of four main subsystems. The receiving element is a pair of pyramidal horn antennas, designed analytically and verified by electromagnetic simulation. At 1.42 GHz, the antennas provide a gain of about 14.2 dBi, with half-power beamwidths of 27.8° and 36.7°. The front-end signal chain uses an LNA → BPF → LNA → HPF configuration and provides approximately 76 dB of gain. The overall noise figure is dominated by the first LNA, with NF ≈ 0.04 dB.

A single USRP B210 software-defined radio serves as the backend of the two-element system. Phase coherence between its two receive channels is maintained by the shared AD9361 oscillator. The digital backend is implemented as an FX correlator in GNU Radio, which performs real-time 8192-point spectral cross-correlation and records complex visibilities for offline analysis.

At the bench level, the system was extended to a four-element, six-baseline configuration using two synchronised USRP B210 units. Phase-coherent operation across all baselines was achieved using a shared 10 MHz reference, a common pulse-per-second timing signal, and a timed-start procedure implemented through the UHD API.

6.1.2 Experimental Validation and Observational Results

The system was validated through solar drift-scan observations at four east–west baselines of 1.00, 1.88, 2.00, and 2.40 m. Interferometric fringes were detected in all four sessions. The

measured fringe periods agreed with the session-averaged Earth-rotation predictions, computed using $\langle \cos H \rangle = 0.974$ to account for the ±1.5 hour observing window around transit, to within 3.4% to 16% across the full range of baselines. The fringe count also increased monotonically with baseline length, rising from 3 peaks at 1.00 m to 9 peaks at 2.40 m over each 3-hour session, in agreement with the expected linear scaling of fringe rate with baseline. Among the four sessions, the observation at $B = 2.00$ m produced the most stable data, with a median visibility magnitude of 31.5 dB and a phase standard deviation of 0.087 rad, the lowest value measured in this study. The visibility amplitude showed no systematic dependence on baseline length, as expected for a source that remains completely unresolved at these separations ($B_{\text{null}} \approx 29.5$ m at 1.42 GHz). The remaining differences between the measured and predicted fringe periods are most plausibly explained by a combination of SNR-related peak selection bias near transit, the ±5 mm uncertainty in baseline measurement, and the absence of formal bandpass calibration.

6.1.3 Key Lessons

This project brought out a few practical points that are not immediately obvious from standard descriptions of interferometry. One of the most important is that phase coherence depends on sharing the same local oscillator, not just a common reference frequency. In the RTL-SDR daisy-chain setup, both devices were driven by the same reference, but their phase-locked loops still developed independent phase offsets. Stable coherence was only obtained when both channels shared a single PLL, as in the AD9361 chip used in the USRP B210.

Extending phase coherence across multiple devices also requires timing control beyond what GNU Radio Companion provides. A PPS-aligned timed-start sequence, implemented at the UHD API level, was essential for achieving repeatable inter-device phase alignment.

Observation site quality also has a strong effect. Moving from a partially obstructed balcony to an open rooftop clearly improved fringe stability and signal-to-noise ratio. The median visibility magnitude improved by 5–6 dB between the balcony and rooftop configurations in comparable sessions. In the early sessions, the main causes of degraded data were obstruction-induced diffraction and RFI reflections.

6.1.4 Limitations

The four-element two-USRP configuration has so far been demonstrated only at the bench level using an injected tone. On-sky fringe observations with this configuration, including calibration to remove the session-dependent LO phase offset, have not yet been carried out.

The current baseline range of 1.00–2.40 m ($u = 4.7$–11.4λ) provides only limited (u, v) coverage. As a result, image reconstruction is not possible with the present data. The instrument can currently measure fringe periods and verify the fringe-rate scaling law, but meaningful

aperture synthesis will require more baselines spanning a wider range of lengths.

Bandpass calibration, which corrects frequency-dependent gain and phase variations across the signal chain, was not applied to the data presented here. This should be done before attempting observations of weaker sources or sources with noticeable spectral structure.

6.2 Future Scope

6.2.1 On-Sky Validation of the Four-Element System

The next immediate step is to test the four-element configuration outdoors and detect solar fringes on all six baselines at the same time. This will require four horn antennas arranged in an east–west line, with each antenna connected through a matched front-end chain to the receive ports of two synchronised USRPs. A calibration signal will also need to be injected at the start of each observing session to measure and correct the inter-device LO phase offset. If successful, this would verify the full operation of the four-element system and provide six measurements of the solar visibility function.

6.2.2 RFSoC-Based Backend

A longer-term improvement to the correlator backend is the use of an RFSoC platform. Unlike the USRP B210, which digitises the signal after analogue downconversion and transfers samples over USB 3.0, an RFSoC such as the Xilinx ZCU216 integrates multi-GHz ADCs and a programmable FPGA on a single chip. This removes the USB bandwidth bottleneck and enables a polyphase filter bank correlator with sidelobe suppression better than $-100\,\mathrm{dB}$, improving on the $-92\,\mathrm{dB}$ achieved by the Blackman–Harris window in the current pipeline [16]. The polyphase filter bank also gives a more uniform channel response across the band, which is important for spectral-line work.

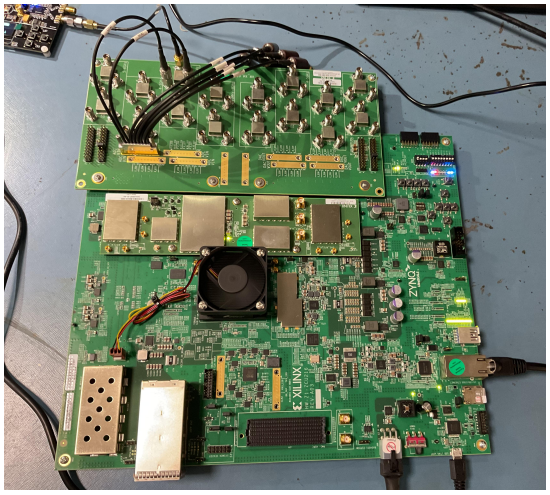

Figure 6.1: AMD Xilinx ZCU216 RFSoC development board used for initial familiarisation and spectrometer development. Preliminary laboratory tests included board bring-up, ADC verification, and integration with the CASPER toolflow.

As part of this project, initial familiarisation with the CASPER toolflow [17], board bring-up of the ZCU216, and a first spectrometer design were completed. A planned continuation of this work is to develop a full multi-element FX correlator on this platform and validate it against the USRP-based system.

6.2.3 Additional Astronomical Targets

Once the system is calibrated, two additional targets are well suited to this instrument. Cassiopeia A and Cygnus A are among the brightest continuum sources in the radio sky at 1.4 GHz. Both are compact enough to appear as point sources at the baselines used here, making them useful for calibration and for fringe tests over a range of hour angles. Detecting fringes from these sources would show that the instrument is sensitive to non-solar targets at flux levels well below that of the Sun.

Galactic HI emission is also accessible in the 1.420 GHz band. At short baselines, the extended emission from the Galactic disc would appear as a slowly varying correlated signal rather than a resolved fringe pattern. Detecting this signal would connect the interferometer more directly to neutral hydrogen astrophysics and also provide a clear reason to explore longer baselines, where finer structure in the HI distribution could be studied.

6.2.4 Aperture Synthesis

The current data set includes only four baselines, all aligned in the east–west direction, which is insufficient for image reconstruction. Extending this toward aperture synthesis would require baselines in at least two orientations. This could be achieved either by rotating the array or by arranging the antennas in a non-collinear configuration. Adopting a non-redundant spacing

would further increase the number of distinct baseline lengths, allowing more spatial frequencies to be sampled simultaneously. With improved (u, v) coverage, it should then be possible to reconstruct a coarse one-dimensional brightness distribution of the Sun using an inverse Fourier transform followed by CLEAN deconvolution [6]. This would represent an initial step toward aperture synthesis with the present instrument.

References

[1] James J. Condon and Scott M. Ransom. *Essential Radio Astronomy*. Princeton Series in Modern Observational Astronomy. Princeton, NJ: Princeton University Press, 2016.

[2] National Radio Astronomy Observatory. *Karl Jansky and his rotating antenna*. Historical image, Accessed: 2026-04-15. URL: https://public.nrao.edu.

[3] Jonathan M. Marr, Ronald L. Snell, and Stanley E. Kurtz. *Fundamentals of Radio Astronomy: Observational Methods*. Series in Astronomy and Astrophysics. Boca Raton, FL: CRC Press, 2016.

[4] National Radio Astronomy Observatory. *Grote Reber's radio telescope, 1937*. Historical image, Accessed: 2026-04-15. URL: https://public.nrao.edu.

[5] National Institute of Standards and Technology. *Robert C. Byrd Green Bank Telescope*. Image credit: NSF/Green Bank Observatory (GBO), Accessed: 2026-04-15. URL: https://www.nist.gov/measuring-cosmos/how-do-stars-and-galaxies-live-and-die%253F/robert-c-byrd-green-bank-telescope.

[6] A. Richard Thompson, James M. Moran, and George W. Swenson. *Interferometry and Synthesis in Radio Astronomy*. 3rd. Astronomy and Astrophysics Library. Open access under CC BY-NC 4.0. Cham: Springer International Publishing, 2017.

[7] National Radio Astronomy Observatory. *Very Large Array (VLA) Technology Page*. Accessed: 2026-04-14. 2024. URL: https://public.nrao.edu/telescopes/vla/technology/.

[8] MIT Haystack Observatory. *The Small Radio Telescope (SRT) for Education*. https://www.haystack.mit.edu/haystack-public-outreach/srt-the-small-radio-telescope-for-education/. Accessed March 2026. 2023.

[9] B. Arul Pandian et al. *Galaxy Rotation Curve Measurements with Low Cost 21 cm Radio Telescope*. arXiv preprint arXiv:2202.11039. Raman Research Institute, Bangalore. 2022.

[10] Jin Koda et al. "A Michelson-type radio interferometer for university education". In: *American Journal of Physics* 84.4 (2016), pp. 249–256.

[11] Timothy Truckle. *Radiation pattern of typical antenna showing main lobe and sidelobes*. https://commons.wikimedia.org/wiki/File:Sidelobes_en.svg. Creative Commons Attribution-ShareAlike License. 2008.

[12] Jayaram N. Chengalur, Yashwant Gupta, and K. S. Dwarkanath, eds. *Low Frequency Radio Astronomy*. 3rd. Available at https://www.ncra.tifr.res.in/ncra/gmrt/gmrt-users/low-frequency-radio-astronomy. Pune, India: National Centre for Radio Astrophysics, TIFR, 2007.

[13] Constantine A. Balanis. *Antenna Theory: Analysis and Design*. 4th. Hoboken, NJ: Wiley, 2016.

[14] David M. Pozar. *Microwave Engineering*. 4th. Hoboken, NJ: Wiley, 2011.

[15] Evan Mayer. *`effex`: Software-defined radio complex correlator (FX architecture, GPU-accelerated)*. GitHub repository, https://github.com/evanmayer/effex. 2021.

[16] Danny C. Price. "Spectrometers and Polyphase Filterbanks in Radio Astronomy". In: *arXiv preprint* arXiv:2106.00558 (2021). Submitted to the Proceedings of URSI GASS 2021.

[17] CASPER Collaboration. *mlib_devel: CASPER FPGA DSP Libraries and Programming Tools*. GitHub repository, https://github.com/casper-astro/mlib_devel. 2024.