



POLITECNICO
MILANO 1863

SCUOLA DI INGEGNERIA INDUSTRIALE
E DELL'INFORMAZIONE

RFI Detection and Suppression in P-Band Spaceborne SAR Data

TESI DI LAUREA MAGISTRALE IN
SPACE ENGINEERING

Author: **Alessandro Bellezza**

Student ID: 10673485

Advisor: Prof. Stefano Tebaldini

Co-advisors:

Academic Year: 2025-26

Abstract

The Biomass mission of the European Space Agency carries the first P-band synthetic aperture radar (SAR) to operate from orbit, with the goal of mapping above-ground forest biomass at global scale. The long P-band wavelength penetrates the forest canopy but confines the radar to a narrow, congested 6 MHz allocation shared with terrestrial radio services. The resulting radio frequency interference (RFI) degrades the calibrated backscatter and the coherence on which the mission's products depend, affecting the weaker cross-polarised channels most: a robust data-driven mitigation method is thus necessary.

This thesis develops, implements, and validates a data-driven RFI detection and suppression chain for Biomass commissioning-phase data. The detector exploits a physical property of quad-polarimetric acquisitions: an external interferer couples identically into the two channels that share a receive antenna port, whereas the scene backscatter in those channels is statistically independent. Their cross-channel coherence, near zero over natural terrain, rises toward unity wherever interference dominates, giving a detection signal that is independent of the interferer's waveform and absolute power and that needs no background-power model. For suppression, a locally estimated Wiener weight uses the cross-polarised channel as an auxiliary reference to coherently cancel the interference while preserving the scene content that spectral notching destroys.

Validated in simulation against an exact ground truth, the detector outperforms the established single-channel reference, most clearly on the weak cross-polarised channels and under heavy contamination. On real Biomass data the coherent cancellation matches the operational ESA filter in image quality while roughly halving the residual cross-channel coherence at the contaminated cells and preserving the natural spectral shape that the operational filter reshapes, showing that the polarimetric structure of the data is itself an effective instrument for removing the interference that corrupts it.

Keywords: Synthetic Aperture Radar; Radio Frequency Interference; Polarimetric coherence; Biomass mission; P-band; Interference cancellation

Abstract in lingua italiana

La missione Biomass dell'Agenzia Spaziale Europea (ESA) impiega il primo radar ad apertura sintetica (SAR) in banda P operante dallo spazio, con l'obiettivo di mappare a scala globale la biomassa forestale epigea. La grande lunghezza d'onda della banda P penetra il volume della chioma forestale, ma costringe il radar a una stretta e trafficata allocazione di 6 MHz condivisa con i servizi radio terrestri. L'interferenza a radiofrequenza (RFI) che ne deriva degrada la retrodiffusione calibrata e la coerenza su cui si fondano i prodotti della missione, colpendo soprattutto i canali cross-polari: una mitigazione robusta e basata sui dati di acquisizione è quindi necessaria.

Questa tesi sviluppa, implementa e valida una catena guidata dai dati per il rilevamento e la soppressione dell'RFI sui dati della fase di commissioning di Biomass. Il rilevatore sfrutta una proprietà fisica delle acquisizioni quad-polarimetriche: un interferente esterno si accoppia in modo identico nei due canali che condividono la stessa porta d'antenna in ricezione, mentre la retrodiffusione della scena in tali canali è statisticamente indipendente. La loro coerenza tra canali (cross-channel), prossima a zero sui terreni naturali, sale verso l'unità ovunque l'interferenza domini, fornendo un segnale di rilevamento indipendente dalla forma d'onda e dalla potenza assoluta dell'interferente. Per la soppressione, un peso di Wiener stimato localmente impiega il canale cross-polare come riferimento ausiliario per cancellare in modo coerente l'interferenza, preservando il segnale che l'azzeramento spettrale (notching) distruggerebbe.

Validato in una simulazione con un'iniezione artificiale di interferenza, il rilevatore supera il metodo di riferimento a singolo canale, specialmente sui deboli canali cross-polari e in condizioni di forte contaminazione. Su dati reali di Biomass la cancellazione coerente è comparabile con il filtro operativo dell'ESA in qualità d'immagine, dimezzando inoltre la coerenza residua tra canali nelle celle contaminate e preservando la forma spettrale naturale che il filtro operativo altera, dimostrando che la struttura polarimetrica dei dati è essa stessa uno strumento efficace per rimuovere l'interferenza.

Parole chiave: Radar ad apertura sintetica; Interferenza a radiofrequenza; Coerenza polarimetrica; Missione Biomass; Banda P; Rimozione dell'interferenza

Contents

Abstract	i
Abstract in lingua italiana	iii
Contents	v
Introduction	1
1 Radar and SAR Signal Processing Fundamentals	5
1.1 Radar Principles and the SAR Concept	5
1.2 Signal Model and Processing	6
1.2.1 Transmitted Waveform and Received Signal	6
1.2.2 Range Compression as Matched Filtering	7
1.2.3 The Range-Frequency Domain and Azimuth Processing	8
1.3 Resolution and the SAR Timing Trade-off	9
1.4 Electromagnetic Scattering and Radiometry	10
1.5 Speckle	11
1.6 Polarimetric SAR	12
1.6.1 Scattering Matrix and Quad-Pol Channels	12
1.6.2 Cross-Channel Coherence	13
1.7 SAR Interferometry and Tomography	13
1.7.1 Interferometric Phase and Coherence	13
1.7.2 SAR Tomography	14
1.8 Linear Estimation and the Wiener Interference Canceller	15
1.8.1 The MMSE Criterion and the Orthogonality Principle	15
2 Radio Frequency Interference in SAR: Sources, Models, and Effects	17
2.1 Sources, Geographical Distribution, and the P-Band Environment	17
2.2 RFI Signal Models	19

2.3	RFI in the Range-Frequency / Azimuth-Time Domain	20
2.4	Effects on SAR Data and Science Products	21
2.5	Detection and Suppression: State of the Art	23
3	The Coherence-Based RFI Detector	27
3.1	Physical Motivation	27
3.2	Working Hypotheses	29
3.2.1	H1 – Scene independence within the receive pair	29
3.2.2	H2 – Common coupling of the interference	29
3.2.3	H3 — Negligible inter-channel acquisition delay	29
3.2.4	H4 – Local interference dominance	30
3.3	The Detection Statistic	31
3.3.1	Coherence under contamination	31
3.3.2	The estimator and its null distribution	31
3.3.3	Two oriented estimators	32
3.4	The Processing Pipeline	33
3.4.1	Defocusing and domain transform	33
3.4.2	Pair-wise coherence and thresholding	33
3.4.3	Line detection	34
3.4.4	Mask application and refocusing	34
3.5	Tuning Parameters	35
3.5.1	Coherence Threshold γ_{th}	35
3.5.2	TSNB Window Length W_F	36
3.5.3	TVWB Window Width W_T	36
3.5.4	Minimum Line Length ℓ_{min}	36
3.6	Comparison with the Meyer Detector	36
3.7	Limitations and Ideal Operating Conditions	37
4	Simulation Environment and Detector Validation	39
4.1	Architecture of the Simulation	39
4.1.1	The defocus–refocus operator pair	40
4.1.2	Synthetic interference injection	41
4.2	Test Scenarios	43
4.3	Performance Metrics	44
4.3.1	Detection metrics	45
4.3.2	Per-channel reconstruction metrics	45
4.3.3	Polarimetric pair-recovery metrics	47
4.4	Results	48

Contents	vii
4.4.1 Detection performance	48
4.4.2 Reconstruction and polarimetric performance	49
4.4.3 Parameter sensitivity and tuning	51
4.5 Discussion	52
5 Application to Real Biomass Data	53
5.1 From Notching to Coherent Cancellation	53
5.1.1 The information loss of notching	53
5.1.2 The cancellation principle	54
5.1.3 Local weight estimation	55
5.1.4 Strengths and weaknesses relative to notching	55
5.2 The Five-Method Comparison	57
5.2.1 Design of the experiment	57
5.2.2 Evaluation without ground truth	58
5.2.3 Results	59
5.3 Interference Environment of the Analyzed Scene	70
5.4 Possible Improvements	71
5.4.1 Azimuth alignment of the receive pair	71
5.5 Discussion on Real Data Results	72
5.5.1 Comparison with the Literature	72
5.5.2 Categorization of the Detected Interferers	74
6 Conclusions and Future Developments	77
6.1 Future Developments	78
6.2 Closing Remark	80
Bibliography	81
A Appendix A	83
List of Figures	89
List of Tables	91
List of Symbols	93
List of Acronyms	97

Introduction

Context and Motivation

Climate change can be considered among the most important and concerning topic that is propelling both scientific and technological progress in this century. Understanding the Global Carbon Cycle and its relation with the increase of CO₂ concentration in the atmosphere is steadily climbing the list of priorities of international organizations, but despite this numerous challenges make obtaining clear and accurate data on the subject almost unfeasible. In particular, Earth's surface role in absorbing atmospheric carbon is well known, and one of the most evident sources of greenhouse-gas concentration reduction every year is situated in broad daylight: it's estimated that Ligneous Vegetation is responsible for storing approximately 45% of all terrestrial carbon, exchanging several billion tonnes of carbon dioxide with the atmosphere every year. Yet the reliability of this estimate of carbon flux, together with the spatial distribution of this stored carbon, quantified as *above-ground biomass* (AGB), remain highly uncertain in the first case and almost unknown at global scale for the second. Ground-based forest inventories are costly, spatially sparse, and difficult to replicate consistently across tropical and boreal regions where access is limited, and the scarcity of data makes model-based estimations and predictions highly uncertain. Logically, space based remote sensing offers the only technically feasible path to global AGB mapping. While space based platforms already performed forestal observation using both optical and microwave wavelengths, no other mission before had the primary scientific objective of measuring Above Ground Biomass: these observations were performed using sensors that were not designed for this purpose, lacking the necessary sensitivity or penetration capability, thus being unable to observe the large woody structures that constitute the bulk of a forest's biomass. [9]

The P-Band (≈ 70 cm wavelength) is the only microwave frequency range with sufficient penetration depth to reach these structures, achieving high biomass sensitivity while also keeping interactions with the ionosphere at a manageable level, making it the only choice for space-based AGB measurement. This scientific imperative motivated the European Space Agency to select the Biomass mission as the seventh Earth Explorer Core Mission,

carrying onboard the first P-Band Synthetic Aperture Radar (SAR) to operate from orbit. All the mission acquisitions will be fully polarimetric, measuring amplitude and phase of the HH, VV, HV and VH channels (For the following part of this document, HV indicates horizontal polarization on transmit and vertical polarization on receive, with the other channels being similarly defined). After an initial commissioning phase, the Tomographic Phase will provide tomographic mapping of all the target areas, across a period of approximately 14 months. For the remaining period of the mission’s lifetime, BIOMASS will carry on the Interferometric Phase, during which 3 passes, each separated by 3 days, will be combined in 3-baseline Pol-InSAR, continuously updating the interferometric maps up until the End of Life (EOL). [9]

The RFI Problem at P-Band

The capabilities and suitability of P-band for this mission come with a clear technical challenge: The SAR allocated Bandwidth for the Biomass mission is restricted by the International Telecommunication Union (ITU) regulations to only 6 MHz around 435 MHz, causing the system’s operating frequency range to be overlapping with a dense population of terrestrial radio services (Fixed and mobile communications networks, digital television broadcasts, military ground-based radars can potentially contribute to the received interference signal. [4]

Table 1: Radar frequency bands for Earth observation.

Band	Frequency	Wavelength	Representative missions
P	0.3–1 GHz	30–100 cm	Biomass (ESA)
L	1–2 GHz	15–30 cm	ALOS-2, NISAR
C	4–8 GHz	3.75–7.5 cm	Sentinel-1, RADARSAT-2
X	8–12 GHz	2.5–3.75 cm	TerraSAR-X, COSMO-SkyMed

This additional RFI power affecting the image would raise the effective noise floor and introduce bright artifacts into the focused SAR images, while at the same time corrupting the backscattered signal’s characteristics that are used for the mission objective of quantifying Biomass and estimating its spatial distribution. Given these concerns, pre-launch airborne campaigns (BioSAR, TropiSAR, AfriSAR) were conducted, among other things, in order to monitor this phenomenon and established that unmitigated RFI can degrade the effective NESZ by 5 to 20 dB in affected areas: without intervention, such biomass retrieval errors would far exceed the mission’s scientific requirements.

The commissioning phase of the mission, during which the satellite transitions from launch to routine science operations, offered the first opportunity to characterize the RFI environment to which the mission will be subject during its remaining lifetime and revealed the presence of RFI corruption with variable intensity across a range of observed geographic areas, including some regions where its occurrence had not been anticipated from pre-launch surveys. Robust, data-driven RFI mitigation algorithms capable of handling such variability are therefore essential. [9]

Objectives and structure of this Thesis

This thesis addresses the development, implementation, and validation of a custom RFI detection and suppression algorithm specifically tailored to the Biomass mission data acquired during its commissioning phase. The work was carried out in collaboration with the Biomass science team at Politecnico di Milano, under the supervision of Prof. Stefano Tebaldini.

In order to achieve the stated goal, the work was divided in two parts: the development of an RFI Detection method which could operate without any prior knowledge about the interferer characteristics (waveform, frequency, bandwidth and location) and that would produce a spatially and spectrally resolved detection mask, exploiting the **cross-channel coherence** between polarimetric channel pairs sharing the same receive antenna port. After being implemented, the Detection step of the algorithm was compared with an established reference method (provided by the paper from Meyer et al. [8]) and validated using a simulated SAR environment with artificially injected RFI. Secondly, two alternative RFI Suppression strategies were evaluated and validated using real Biomass commissioning-phase data, acquired across areas with different RFI intensity levels, and as a final step these removal methods were compared with the official ESA Biomass algorithm currently deployed.

After a Radar and SAR imaging fundamentals explanation [1], where all concepts used in the development of the algorithm are presented, an RFI overview is carried out in [2]: both its physical origins and analytical models are analyzed in order to present its defining characteristics, which could be used as targets for various detection methods. Consequently, the implemented RFI Detector algorithm is described in [3], together with its own necessary assumptions, known limitations and ideal operating conditions. The need for validating the proposed detection method, as well as evaluating its performances in a controlled and repeatable environment made it necessary to implement a simulation where clean SAR images could be corrupted with known artificial interferences and

subsequently given as input to each analyzed method, described in [4]. The simulated RFI benchmark results, followed by the application of the same detector to real Biomass corrupted data, are presented in [5].

1 | Radar and SAR Signal Processing Fundamentals

1.1. Radar Principles and the SAR Concept

A radar (RAdio Detection And Ranging) system is an active sensor: it transmits a radio-frequency pulse and records the echo backscattered by the scene. The range to a scatterer follows from the round-trip propagation delay Δt as $R = c \Delta t / 2$, where c is the speed of light. Because the phase of the transmitted waveform is known, the received echo preserves both amplitude and phase, meaning the system is coherent. This allows phase differences between images to be interpreted as surface displacements or height differences, and it is what could make interference from external radio sources detectable through cross-channel correlation.

A real-aperture radar achieves azimuth resolution $\delta_x = \lambda R / L_x$, which degrades linearly with range R : at spaceborne ranges of 600–700 km and P-Band wavelength, achieving a resolution in the order of tens of meters would require an antenna several tens of kilometers long. The Synthetic Aperture Radar (whose concept was first developed by mathematician Carl A. Wiley in 1951) avoids this issue using a compact physical antenna, recording coherent echoes at each position along the orbit, and then processing them jointly as if they were signals collected by a much longer discrete array. The key to this process is that the phase of each echo is preserved: because $\exp(-j4\pi R(\tau)/\lambda)$ varies deterministically with the platform position τ , the contributions from all pulses illuminating a given ground point can be phase-aligned and summed coherently. The resulting azimuth resolution is independent of range and wavelength, obtaining a constant resolution value of $\delta_x = L_x/2$ when the full synthetic aperture is processed.

For the remaining of this work, the following reference system convention have been used: the SAR standard coordinate system places the x axis along the flight direction (azimuth), the y axis across track in the ground plane (ground range), the z axis vertically (elevation), and uses slant range r along the line-of-sight at incidence angle θ . The relationship

between slant-range and ground-range resolution is $\delta_y = \delta_r / \sin \theta$. All systems considered here are monostatic and operate in strip-map mode, where the antenna footprint sweeps a continuous swath as the platform advances.

1.2. Signal Model and Processing

1.2.1. Transmitted Waveform and Received Signal

SAR systems transmit a linear frequency-modulated (LFM) pulse, also known in Radar jargon as Chirp, whose complex envelope is a rectangular wave amplitude function, with a linearly increasing instantaneous frequency:

$$g(t) = \text{rect}\left(\frac{t}{T_p}\right) e^{(j\pi K_r t^2)} \quad (1.1)$$

where T_p is the pulse duration and K_r is the chirp rate in Hz/s. The instantaneous frequency sweeps linearly over the pulse duration, for a total bandwidth $B = K_r T_p$ Hz. The LFM waveform decouples pulse duration (which controls the transmitted energy $E = P_{\text{tx}} T_p$) from range resolution, which is set entirely by B . This reason explain why the chirp pulse have become the standard in all Radar systems: long, low-peak-power pulses can therefore yield fine range resolution, allowing SAR systems to meet sensitivity requirements without exceeding hardware power limits.

After modulation onto carrier frequency f_0 , propagation to a point scatterer at range $R(\tau)$ and back, and demodulation at the receiver, the received signal (baseband) from that scatterer is:

$$s(t, \tau) = A_t g\left(t - \frac{2R(\tau)}{c}\right) e^{(-j\frac{4\pi}{\lambda} R(\tau))} \quad (1.2)$$

where t is fast time (time within a pulse), τ is slow time (pulse emission time along the trajectory), $\lambda = c/f_0$ is the wavelength, and A_t is the complex reflectivity of the scatterer. It's common practice in aerial and spaceborne SAR, due to the ratio between pulse propagation velocity c and platform movement speed v_{SAR} , to assume that transmission and reception happen at the same time instant relative to the slow time τ (Start-Stop approximation). The exponent $(-j4\pi R/\lambda)$ encodes the two-way propagation phase: it changes by 2π for every wavelength of range variation, both for the incident and backscattered pulse, resulting in a total of $4\pi R$. For a distributed scene the received signal is the coherent superposition of contributions from all scatterers within the illuminated area, each with its own A_t and $R(\tau)$.

Thermal noise $n(t, \tau)$ – commonly modeled as a complex circular Gaussian process with zero mean and one-sided spectral density $N_0 = K_B T_{\text{sys}}$, where K_B is Boltzmann's constant and T_{sys} is the effective noise temperature – is superimposed on the scene return. Furthermore, since Biomass operating band is shared with terrestrial radio services, an additional interference term $r(t, \tau)$ may be present. The baseband data then take the form:

$$h(t, \tau) = s(t, \tau) + r(t, \tau) + n(t, \tau) \quad (1.3)$$

The collection of received pulses, stored as a 2-D matrix indexed by (t, τ) (meaning that each received pulse is stored as a column whose index is the transmission/reception time) is known as the raw data matrix $D(t, \tau)$. [12]

1.2.2. Range Compression as Matched Filtering

The first processing step, range compression, maximises the SNR at the output of the range detector and concentrates the energy of each echo into a narrow peak along the t axis. Its theoretical basis is the matched filter: consider the problem of estimating the delay of a known signal $g(t)$ received in additive white Gaussian noise: $h(t) = A g(t - \tau_0) + n(t)$. Among all linear filters $u(t)$, the one that maximises the output signal-to-noise ratio at time τ_0 is the time-reversed conjugate replica of the signal, $u(t) = g^*(-t)$. The proof follows directly from the Schwarz inequality:

$$\left| \int h(t) u(t) dt \right|^2 \leq \int |h(t)|^2 dt \cdot \int |u(t)|^2 dt \quad (1.4)$$

with equality when $u(t) \propto h^*(t)$. Applied to the baseband received signal, the matched filter output, or range compressed pulse (d_{rc}) at a given slow-time position τ is the cross-correlation of the received pulse with a clean replica of the transmitted waveform:

$$d_{\text{rc}}(t, \tau) = \int g_{TX}^*(\xi - t) g_{RX}(\xi, \tau) d\xi = (g_{RX} \star g_{TX})(t, \tau) \quad (1.5)$$

where $\star$ denotes a linear convolution, and the result of this operation is the cross-correlation between the two convoluted signals. For a point target at delay $t_0 = 2R/c$, this evaluates to:

$$d_{\text{rc}}(t, \tau) = A_t T_p \text{sinc}[B(t - t_0)] \exp\left(-j \frac{4\pi}{\lambda} R(\tau)\right) \quad (1.6)$$

The cardinal sine envelope has a 3-dB width of $1/B$ (the energy accumulated across the pulse duration is concentrated into a peak of width $\approx 1/B$), corresponding to a slant-

range resolution of $\delta_r = c/(2B)$. For Biomass with $B = 6$ MHz, this gives $\delta_r \approx 25$ m. [2, 12]

After this definition, it's clear that range compression is fundamentally a correlation operation: it measures the similarity between the received data and a delayed copy of the transmitted waveform. This perspective will recur in azimuth compression and, later, in the definition of interferometric coherence.

1.2.3. The Range-Frequency Domain and Azimuth Processing

Range compression produces a 2-D data matrix $D_{\text{rc}}(t, \tau)$ in which the range direction has been compressed to resolution δ_r , but the slow-time axis retains the full phase history of every scatterer.

As the platform moves along the orbit with velocity v , the slant range to a target at broadside range R_0 varies according to:

$$R(\tau) = \sqrt{R_0^2 + v^2\tau^2} \approx R_0 + \frac{v^2\tau^2}{2R_0} \quad (1.7)$$

This parabolic range history, called range migration, causes the target signature to spread across fast-time positions in the raw data matrix. Azimuth compression must simultaneously correct for this migration and integrate the phase history coherently.

The azimuth processing can be understood as a second matched filtering, now along the slow-time axis. [12]

An alternative focusing method that does not rely on the same approximations as range and azimuth matched filtering, at the cost of increased computational cost is Time-Domain Back-Projection (TDBP). For each output pixel (x_i, r_i) , TDBP evaluates the range-compressed data at the correct delay for every pulse, removes the phase $\exp(-j4\pi R/\lambda)$, and sums:

$$I(x_i, r_i) = \int d_{\text{rc}}\left(\frac{2R(\tau; x_i, r_i)}{c}, \tau\right) \exp\left(j\frac{4\pi}{\lambda}R(\tau; x_i, r_i)\right) d\tau \quad (1.8)$$

The integral sums coherently only where the applied phase exactly compensates the propagation phase, i.e. at the true target position. At all other positions the contributions add with random phases and cancel statistically. TDBP requires no approximation about the trajectory shape or the signal bandwidth and is therefore exact; it also handles arbitrary platform motion, making it the standard focuser for airborne and space-based systems where orbital perturbations are significant. As mentioned before, the computational cost scales as $O(N_r N_a^2)$ for a scene of $N_r \times N_a$ pixels, which is the principal practical drawback

relative to frequency-domain algorithms.

independently of the focusing method used, the focused complex image for a single point target takes the form of a 2-D sinc function:

$$I(x, r) = A_t \operatorname{sinc}\left(\frac{x - x_t}{\delta_x}\right) \operatorname{sinc}\left(\frac{r - r_t}{\delta_r}\right) \exp\left(-j \frac{4\pi}{\lambda} R_0\right) \quad (1.9)$$

where $\delta_x = L_x/2$ is the azimuth resolution, R_0 is the broadside range, and the exponential retains the absolute propagation phase that is exploited in interferometry. For a distributed scene the focused image is the 2-D convolution of the scene's complex reflectivity distribution with this impulse response. [2]

1.3. Resolution and the SAR Timing Trade-off

The three primary resolution parameters of a SAR system follow directly from the signal model. Slant-range resolution is $\delta_r = c/(2B)$, set exclusively by the transmitted bandwidth when exploiting the properties of transmitting LFM pulses. Ground-range resolution projects this onto the horizontal plane: $\delta_y = \delta_r / \sin \theta$, degrading at shallow incidence angles and providing the most stringent requirements on observation geometry, and thus orbit choice. Similarly Azimuth resolution (when the full physical aperture is processed) is driven by an independent system parameter and is $\delta_x = L_x/2$, fully decoupling the design choices to achieve the desired resolutions.

the final system design parameter is the Pulse Repetition Interval (PRI), whose choice involves a compromise: to avoid range ambiguities (receiving an echo from one pulse simultaneously with the echo of the previous pulse) the PRI must exceed the round-trip delay difference across the swath: $\text{PRI} > 2(r_{\max} - r_{\min})/c$

At the same time, the along-track sampling step $\Delta x = v \cdot \text{PRI}$ must satisfy the Nyquist criterion for the azimuth spectrum, requiring $\text{PRI} < \lambda/(4v)$ when the full aperture is processed. This gives the allowed PRI interval to avoid range and azimuth ambiguities:

$$2(r_{\max} - r_{\min})/c < \text{PRI} < \lambda/(4v) \quad (1.10)$$

A wider swath tightens the lower bound on the PRI, while fine azimuth sampling tightens the upper bound: the two requirements compete, and neither can be satisfied arbitrarily without trading off the other.

A resolution cell of area $\delta_x \times \delta_y$ defines the spatial support of a single image pixel. All

scatterers within one cell contribute coherently to that pixel's complex value, which is a random process when observing a distributed target due to the random variations in the contributions of the single elements in the same resolution cell. This issue is at the origin of the phenomenon known as Speckle in SAR images, where different observations in space or time present random magnitude variations across the image spatial coordinates (Explored more in depth in [1.5]). To reduce this effect, a technique known as Multilooking is commonly applied: By averaging L_{eq} statistically independent looks (averaging over spatially adjacent cells) the variance of the pixel's magnitude is reduced at the cost of reduced spatial resolution. [12]

1.4. Electromagnetic Scattering and Radiometry

The amplitude of the focused SAR image reflects the strength of the backscattered field from the scene. For a distributed target occupying one resolution cell, the radar cross section (RCS) is $\sigma = \sigma^0 \delta_x \delta_y$, where σ^0 (normalised RCS, or backscatter coefficient) is an intrinsic property of the surface material, roughness, and viewing geometry. The received power from that cell, at range R and incidence angle θ , follows from the radar equation:

$$P_{\text{rx}} = \frac{P_{\text{tx}} G^2 \lambda^2 \sigma^0 \delta_x \delta_y}{(4\pi)^3 R^4 \sin \theta} \quad (1.11)$$

where G is the antenna gain (equal for transmit and receive in the monostatic case). The post-compression SNR, after integrating N_τ pulses within the synthetic aperture, is $\text{SNR} = P_{\text{rx}} T_p N_\tau / (K_B T_{\text{sys}})$.

In SAR imaging, a parameter equivalent to the SNR is commonly used, known as the Noise-Equivalent Sigma Zero (NESZ) (defined as the backscatter level at which SNR equals unity):

$$\sigma_{\text{NESZ}}^0(\theta) = \frac{(4\pi)^3 R^4(\theta) \sin \theta K_B T_{\text{sys}}}{P_{\text{tx}} G^2 \lambda^2 T_p N_\tau} \quad (1.12)$$

Where the relation between NESZ and SNR is $\text{SNR} = \sigma^0 / \sigma_{\text{NESZ}}^0$. ESA's Biomass mission requires $\text{NESZ} \leq -27$ dB across the full swath, a stringent constraint that motivates the large 12-m deployable reflector antenna, and the necessity for interference mitigation when localized high power interferers degrade the system performances below the required threshold. [2, 12]

1.5. Speckle

A natural-terrain resolution cell contains many elementary scatterers whose positions vary randomly at the scale of λ . The complex pixel value is the coherent superposition of their individual contributions:

$$u = \sum_k A_k e^{j\phi_k} \quad (1.13)$$

When the number of scatterers is large and none dominates, the central limit theorem applies and u converges to a complex circular Gaussian (CCG) random variable with zero mean. Its real and imaginary parts (defined as the in-phase (I) and quadrature (Q) components) are independent Gaussian variables of equal variance $\sigma_u^2/2$, giving an intensity $|u|^2$ that is exponentially distributed with mean $\sigma_u^2 = \mathbb{E}[|u|^2]$ proportional to σ^0 . This characteristic random variation of intensity in the image, called speckle, is a consequence of the coherent nature of the illumination and is inevitable within a single image.

Multi-looking reduces the speckle variance by averaging L_{eq} independent intensity values, so that the averaged intensity converges toward its mean σ^0 with a variance that decreases as $(\sigma^0)^2/L_{\text{eq}}$. The choice of L_{eq} is once again a trade-off opportunity between spatial resolution and the precision of any statistical estimate.

The CCG model also determines the statistics of the cross-correlation between two images. If u_1 and u_2 are CCG with zero mean, their (population) complex cross-correlation is $\mathbb{E}[u_1 u_2^*]$, and the *coherence* between them is:

$$\gamma_{12} = \frac{\mathbb{E}[u_1 u_2^*]}{\sqrt{\mathbb{E}[|u_1|^2] \mathbb{E}[|u_2|^2]}} \quad (1.14)$$

with $|\gamma_{12}| \in [0, 1]$. When u_1 and u_2 are statistically independent (as is the case for co-pol and cross-pol channels over terrain with negligible slope) the numerator vanishes and $|\gamma_{12}| = 0$. When the two signals share a common component, the coherence magnitude increases proportionally to that component's fractional power. This is the direct connection between cross-channel correlation and the presence of a shared source, which the custom RFI detector will exploit.

In practice, the population coherence is estimated from a finite window of ℓ_{coh} independent looks:

$$\hat{\gamma}_{12} = \frac{\sum_{\ell} u_1^{(\ell)} \left(u_2^{(\ell)}\right)^*}{\sqrt{\sum_{\ell} |u_1^{(\ell)}|^2 \cdot \sum_{\ell} |u_2^{(\ell)}|^2}} \quad (1.15)$$

Under the ideal hypothesis with $\gamma_{12} = 0$ (uncorrelated clean channels), the expected magnitude of the coherence estimate $\hat{\gamma}$ decreases as $\sim 1/\sqrt{L_{\text{eq}}}$, so larger windows yield more accurate coherence estimation and lower the chance of falsely categorizing a pixel as RFI affected due to a spurious coherence spike. [12]

1.6. Polarimetric SAR

1.6.1. Scattering Matrix and Quad-Pol Channels

The polarisation of an electromagnetic wave describes the orientation of its electric field vector. Decomposing it into horizontal (H) and vertical (V) components relative to the ground plane, the monostatic backscatter relationship between transmitted and received polarisation states is given by the 2×2 complex scattering matrix:

$$\begin{bmatrix} E_H^r \\ E_V^r \end{bmatrix} = \frac{e^{-j2kR}}{R} \begin{bmatrix} S_{HH} & S_{HV} \\ S_{VH} & S_{VV} \end{bmatrix} \begin{bmatrix} E_H^t \\ E_V^t \end{bmatrix} \quad (1.16)$$

where subscripts denote transmit and receive polarisation in that order: S_{HV} is received on H when V is transmitted. A fully polarimetric SAR alternates H and V transmissions on successive pulses while simultaneously receiving on both H and V, recovering all four elements of $\mathbf{S}$. It has to be noted that the four channels are not independent, instead for the reciprocity theorem, which holds for non-moving targets, $S_{HV} = S_{VH}$ so there are 3 independent components for each pixel's scattering matrix. [5]

The scattering information can be also expressed by the 3×3 polarimetric covariance matrix formed from the Pauli scattering vector $\mathbf{k} = [S_{HH}, \sqrt{2}S_{HV}, S_{VV}]^T$:

$$\mathbf{C} = \langle \mathbf{k} \mathbf{k}^H \rangle \quad (1.17)$$

where angle brackets denote spatial averaging. The diagonal entries are the channel backscatter powers; the off-diagonal entries encode the physical cross-correlations between channels.

The four focused images I_{HH} , I_{HV} , I_{VH} , I_{VV} are then called the polarimetric channels. Two of them share the H receive port (I_{HH} and I_{HV}) and two share the V receive port (I_{VV} and I_{VH}). This shared-antenna architecture has a direct consequence for any external signal entering the system: it arrives as a physical wavefront with a fixed polarimetric signature and is distributed to both channels served by the same receive port, with a delay defined by the time interval between the H-transmitted pulse and the V transmitted pulse

emission. A signal that is external to the scene and is emitted as H-polarized therefore couples identically (up to a phase difference determined by the pulse repetition interval) into both I_{HH} and I_{HV} , and similarly for a V-polarized signal into I_{VV} and I_{VH} . [5]

1.6.2. Cross-Channel Coherence

Over natural distributed terrain with no slope, the ensemble cross-correlation between co-pol and cross-pol channels vanishes: $\mathbb{E}[S_{HH}S_{HV}^*] = 0$. This is a consequence of the scattering symmetry: the mechanisms that produce co-pol returns (specular surfaces, even-bounce structures) are physically distinct from those producing cross-pol returns (volume scattering, odd-order multiple bounces with orientation variation), and their contributions have to be statistically uncorrelated. The coherence $|\gamma_{HH,HV}|$ estimated from clean data over such terrain is therefore near zero.

The situation changes when an external signal is present. An interferer entering the H receive chain contributes the same waveform $r_H(f, \tau)$ to both I_{HH} and I_{HV} . In the range-frequency / azimuth-time domain the total data in the H-receive pair is:

$$H_{HH}(f, \tau) = S_{HH}(f, \tau) + r_H(f, \tau) + n_{HH}(f, \tau) \quad (1.18)$$

$$H_{HV}(f, \tau) = S_{HV}(f, \tau) + r_H(f, \tau) + n_{HV}(f, \tau) \quad (1.19)$$

When the interference power dominates at a particular (f, τ) position ($|r_H|^2 \gg |S_{HH}|^2, |S_{HV}|^2$) the cross-correlation $\mathbb{E}[H_{HH}H_{HV}^*] \approx \mathbb{E}[|r_H|^2]$, driving $|\hat{\gamma}_{HH,HV}|$ toward unity. The cross-channel coherence therefore transitions sharply from near zero in clean regions to near one wherever interference dominates, regardless of the interferer's waveform or frequency. The same argument applies to the V-receive pair (I_{VV}, I_{VH}) independently. [5]

1.7. SAR Interferometry and Tomography

1.7.1. Interferometric Phase and Coherence

When two coherent SAR images I_1 and I_2 of the same scene are acquired from slightly different positions or times, their pixel-wise complex product – known as interferogram – ideally carries a phase proportional to the range difference:

$$\phi_{21} = \angle(I_2 I_1^*) = \frac{4\pi}{\lambda}(R_1 - R_2) \quad (1.20)$$

More precisely, the interferometric phase decomposes as:

$$\phi_{21} = k_z \delta z + \frac{4\pi}{\lambda} d_{\text{LoS}} + \phi_{\text{APS}} + \phi_{\text{noise}} \quad (1.21)$$

where $k_z = 4\pi b_{\perp}/(\lambda R \sin \theta)$ is the height-to-phase conversion factor ($b_{\perp}$ being the perpendicular orbital baseline), d_{LoS} is the surface displacement along the line of sight, and ϕ_{APS} is the differential atmospheric delay. The term ϕ_{noise} represents the phase error arising from incomplete coherence between the two acquisitions. [1]

The interferometric coherence γ (as defined in Equation (1.14), with $u_1 = I_1$ and $u_2 = I_2$) summarises all mechanisms that reduce the correlation between the images. Its magnitude decomposes multiplicatively:

$$|\gamma| = |\gamma_{\text{SNR}}| \cdot |\gamma_{\text{temporal}}| \cdot |\gamma_{\text{geometric}}| \cdot |\gamma_{\text{volume}}| \quad (1.22)$$

where $|\gamma_{\text{SNR}}| = 1/(1 + \sigma_{\text{NESZ}}^0/\sigma^0)$ is thermal noise decorrelation, $|\gamma_{\text{temporal}}|$ captures temporal scene change, $|\gamma_{\text{geometric}}|$ accounts for baseline decorrelation (spectral shift between the two acquisition geometries), and $|\gamma_{\text{volume}}|$ reflects scattering within a volumetric target. In the presence of RFI that is incoherent between acquisitions, a fifth term $|\gamma_{\text{RFI}}|$ further influences the coherence in contaminated pixels.[1]

The phase error variance is bounded below by the Cramér-Rao bound:

$$\sigma_{\phi}^2 \geq \frac{1 - |\gamma|^2}{2L_{\text{eq}}|\gamma|^2} \quad (1.23)$$

This expression shows that phase noise increases non-linearly as coherence decreases: even a modest reduction from, say, $|\gamma| = 0.9$ to $|\gamma| = 0.6$ roughly triples the phase variance. The coherence is therefore the primary quality indicator for interferometric and tomographic products, and any processing step – including RFI suppression – must be evaluated for its effect on it. [12]

1.7.2. SAR Tomography

SAR Tomography extends the interferometric concept to three dimensions by processing N repeat-pass acquisitions at N distinct orbital baselines $\{b_n\}$ jointly, in order to expand the 2D resolution cell concept to 3D, adding an elevation resolution component to the slant-range and azimuth. The n -th focused image at a given (x, r) pixel contains the

integral:

$$I_n = \int s(v) \exp(-jk_{z,n}v) dv, \quad k_{z,n} = \frac{4\pi b_n}{\lambda R_0 \sin \theta} \quad (1.24)$$

where $s(v)$ is the complex reflectivity profile along the elevation coordinate v and $k_{z,n}$ is the spatial frequency in elevation sampled by the n -th pass. The data vector $\mathbf{I} = [I_1, \dots, I_N]^T$ contains N discrete Fourier samples of $s(v)$; tomographic inversion recovers the full vertical reflectivity profile. The achievable elevation resolution is $\delta_v = \lambda R_0 \sin \theta / (2\Delta b)$, where Δb is the total baseline span. For the Biomass tomographic acquisition phase, the orbital tube constrains Δb to a few hundred metres, which at P-Band yields elevation resolutions in the tens of metres: enough to separate canopy, mid-canopy, and ground scattering layers in forests. [12]

1.8. Linear Estimation and the Wiener Interference Cancellation

1.8.1. The MMSE Criterion and the Orthogonality Principle

A recurring problem in signal processing is the estimation of a signal of interest s from a noisy or contaminated observation. When the estimator is constrained to be a linear function of the observations, the minimum mean-square error (MMSE) criterion selects the filter that minimises $\mathbb{E}[|s - \hat{s}|^2]$, where $\hat{s} = \int w(\xi) y(\xi) d\xi$ is the linear estimate of s from observations y .

In the frequency domain, for stationary processes, the MMSE criterion reduces to the Wiener-Hopf equation: the optimal filter spectral response is

$$W_{\text{opt}}(f) = \frac{S_{ys}(f)}{S_{yy}(f)} \quad (1.25)$$

where $S_{ys}(f) = \mathbb{E}[\tilde{y}(f)\tilde{s}^*(f)]$ is the cross-power spectrum between the observation and the desired signal, and $S_{yy}(f) = \mathbb{E}[|\tilde{y}(f)|^2]$ is the auto-power spectrum of the observation. The minimum achievable MSE is $\mathbb{E}[|s|^2] - \mathbb{E}[|W_{\text{opt}}\tilde{y}|^2]$, which decreases as the correlation between y and s increases. [3]

2 | Radio Frequency Interference in SAR: Sources, Models, and Effects

The frequency range below 1 GHz that makes long-wavelength SAR scientifically distinctive is also densely occupied by terrestrial radio services. Any spaceborne receiver operating in this range necessarily collects electromagnetic energy from transmitters on the ground whose waveforms were not designed to interact with a radar sensor and whose presence degrades the SAR measurement in ways that depend on both the interferer's spectral character and the signal-processing domain in which the degradation manifests. This chapter treats the physical and mathematical aspects of RFI as they are relevant to the detection and suppression methodology developed in this work. The development proceeds from the sources and models of interference, through its representation in the processing domain where it is most clearly separated from the SAR signal, to its quantitative effect on the image and coherence products that constitute the mission's science output. [10]

2.1. Sources, Geographical Distribution, and the P-Band Environment

The International Telecommunication Union partitions the radio spectrum between satellite services and terrestrial services through a treaty-based allocation system. For Earth observation SARs operating at P-Band, the relevant allocation is a 6 MHz window centred on 435 MHz within the 420–470 MHz band, granted to the Earth Exploration Satellite Service (EESS) on a *secondary* basis. Primary status belongs to fixed, mobile, radiolocation, and amateur services, which means that spaceborne sensors must tolerate interference from these services and have no legal mechanism to demand their silence [4]. The Biomass mission operates within this framework: it uses the full allocated 6 MHz and accepts that in-band RFI is a permanent feature of its data.

The terrestrial transmitters that contribute to the P-Band interference environment fall into three broad physical categories. *Communication services* — including TETRA professional mobile radio networks, analogue legacy networks in transition toward digital, and military tactical radio links — generate narrowband, frequency-stable carriers that remain active continuously or quasi-continuously. *Broadcasting services* contribute through harmonic and inter-modulation products of higher-power AM and FM stations whose fundamental frequencies are below 435 MHz, reaching the P-Band allocation through non-linear propagation effects. *Radar services* — air-traffic control primary radars, maritime vessel traffic systems, and military surveillance radars — produce pulsed, wideband emissions that are intermittent and potentially broadband relative to the SAR receiver window. [10]

Pre-launch characterisation of the P-Band interference environment, conducted through dedicated airborne campaigns and from existing L-Band SAR datasets repurposed as RFI monitors, found that Europe and the Middle East exhibit the densest narrowband interference, attributable primarily to TETRA networks [8]. Central and East Asia show intermittent wideband contamination from military radars. The Amazon basin, Congo basin, and other tropical forest regions — the primary science targets of Biomass — generally carry lower interference power, but dedicated in-flight surveys prior to the mission’s first science phase revealed that interference was present in some of these areas at levels that had not been predicted from ground-based spectrum measurements. This geographic variability, and the presence of interference at unexpected locations, motivated the development of a data-driven detection algorithm that does not rely on a pre-compiled interference catalogue. [10]

The physical mechanism by which a terrestrial transmitter contaminates a SAR acquisition is straightforward. The spaceborne antenna illuminates a ground swath and, during reception, is sensitive to all electromagnetic energy arriving from the corresponding solid angle. A transmitter within the antenna footprint – or outside it but within the far side-lobes of the large Biomass 12-m reflector – will cause its signal to couple into the receiver. The received interference power follows a path-loss law distinct from the SAR backscatter (which travels twice the slant range): because the interferer radiates isotropically or with limited directivity from the ground, the interference arrives only once at the antenna, reducing the range-dependent attenuation relative to the two-way SAR signal. At P-Band with its large wavelength and wide lobes, strong interferers can be detected even when they lie tens of kilometres outside the nominal swath boundary.[12]

2.2. RFI Signal Models

The baseband received signal, after demodulation and prior to any filtering, is $h(t, \tau) = s(t, \tau) + r(t, \tau) + n(t, \tau)$ as introduced in Equation (1.3). The interference term $r(t, \tau)$ can be characterised by its behaviour along both the fast-time (t) and slow-time (τ) axes, which together determine where and how it manifests in the processing chain.

Time-Stationary Narrowband interference (TSNB). A narrowband interferer – a communication carrier, a CW beacon, a stable radar harmonic – produces a baseband tone at a fixed frequency ν_0 within the SAR receiver bandwidth. Its general form is:

$$r_{\text{NB}}(t, \tau) = A_0(\tau) \exp\left(j2\pi\nu_0 \frac{t}{f_s} + j\psi(\tau)\right) \cdot L_{\text{az}}(\tau) \quad (2.1)$$

where f_s is the sampling rate, $A_0(\tau)$ is a slowly varying amplitude, $\psi(\tau)$ is a slowly varying phase unrelated to the radar chirp, and $L_{\text{az}}(\tau) \in \{0, 1\}$ indicates which slow-time positions receive the interference. When the interferer is continuously active, $L_{\text{az}} = 1$ for all τ and A_0 is approximately constant, producing what the literature terms a Time-Stationary Narrowband (TSNB) signal [8]. If the interferer or its propagation channel is intermittent, L_{az} takes a random or periodic binary pattern.

The interference bandwidth is determined by the modulation state of the source. A pure CW carrier occupies a single range-frequency bin; narrowband modulated signals (e.g. TETRA with a channel bandwidth of 25 kHz, much smaller than the 6-MHz SAR bandwidth) occupy a contiguous block of a few bins. Because a sharp gate in time would produce sinc-shaped sidelobes in frequency, a tapered range envelope is often assumed when modelling this kind of communication signals.

Time-Varying Wideband interference (TVWB). Pulsed emitters – ATC primary radars, maritime VTS systems, military surveillance radars – transmit wideband LFM or noise-modulated pulses at a sub-repetition rate that is generally incommensurate with the SAR PRF. The SAR receiver window may capture part of one burst, all of it, or an overlap between consecutive bursts. The general TVWB model is:

$$r_{\text{WB}}(t, \tau) = A_0 \cdot w_r(t) \cdot s_{\text{wb}}(t) \cdot L_{\text{az}}(\tau) \quad (2.2)$$

where $w_r(t)$ is the burst gate envelope (non-zero only within the burst duration T_d), $s_{\text{wb}}(t)$ is the waveform content (LFM, noise, or a sub-pulse train), and $L_{\text{az}}(\tau)$ selects the slow-time positions where the burst falls within the receive window. For a deterministic

wideband chirp the waveform is:

$$s_{\text{wb}}(t) = \exp\left(j2\pi\left[f_c t + \frac{K_{\text{wb}}}{2}t^2\right]\right), \quad K_{\text{wb}} = \frac{B_{\text{wb}}}{T_d} \quad (2.3)$$

sweeping bandwidth B_{wb} over the burst duration. For a stochastic noise emitter, $s_{\text{wb}}(t)$ is a complex CCG process bandlimited to $[f_c - B_{\text{wb}}/2, f_c + B_{\text{wb}}/2]$. A practical third variant is a train of sub-pulses at sub-repetition rate PRF_{sub} , each occupying duration τ_p and bandwidth B_{sub} , modelling fast-scanning radar systems whose inter-pulse period is shorter than the SAR receive window.

In all variants the Signal-to-Interference Ratio (SIR), defined as $\text{SIR}_{\text{dB}} = 10 \log_{10}(P_s/P_r)$ where P_s is the mean SAR signal power and P_r is the mean interference power within the affected spectral region, parametrizes the contamination severity. Observations at P-Band from airborne campaigns indicate that SIR values from -20 dB (interference power $100\times$ the scene return) to $+10$ dB (interference comparable to strong backscatter) are all encountered in practice, depending on the distance to the terrestrial emitter and the scene's backscatter level [8].

2.3. RFI in the Range-Frequency / Azimuth-Time Domain

The processing step that most clearly separates RFI from the SAR signal is the transformation to the range-frequency / azimuth-time domain (f, τ) defined by taking the Fourier transform of each baseband pulse along the fast-time axis. In this representation the two signal classes occupy distinct regions: the SAR signal, after chirp inversion, has a flat power spectrum spanning the full bandwidth $[-B/2, B/2]$ with a spectral density approximately equal to $\sigma_{\text{avg}}^0 \cdot \delta_r/P_{\text{rx}}$ per bin. Its phase varies rapidly with both f (the residual chirp phase after range compression) and τ (the Doppler phase history of each scatterer), and there is no systematic correlation between the range spectra of successive azimuth lines.

TSNB interference, by contrast, concentrates at isolated frequency positions ν_0 and maintains that concentration across all azimuth lines where the interferer is active. In the (f, τ) domain it appears as a bright horizontal stripe: a narrow vertical extent in frequency (determined by the interference bandwidth and the range-gate taper) combined with full or partial azimuth extent (determined by L_{az}). The spectral power at $f = \nu_0$ may exceed the SAR broadband level by 20–40 dB for strong narrowband sources.

TVWB interference produces the complementary morphology: a bright vertical stripe confined to the range-frequency positions $\tilde{w}_r(f) \cdot \tilde{s}_{wb}(f)$ and to the azimuth positions $L_{az}(\tau)$ where the burst falls within the receive window. For a broadband LFM burst, the contaminated frequency region spans $[f_c - B_{wb}/2, f_c + B_{wb}/2]$ uniformly, while for noise interference it is a flat spectral floor elevated above the SAR background.

This morphological separation – TSNB as horizontal features, TVWB as vertical features in the (f, τ) plane – is the reason the detection algorithm (Chapter 3) uses two complementary sliding-window coherence estimators that integrate along different axes. The orthogonality of the two signatures also means that the two types of interference are relatively easy to separate from each other and to address with distinct processing strategies, even when both are present simultaneously.

Two statistical properties of the interference distinguish it from the thermal noise background. First, the inter-pulse correlation: for a persistent narrowband source, the complex spectrum $\tilde{r}(f_0, \tau)$ at the interference frequency is a slowly varying function of τ , while the SAR clutter at any fixed frequency is essentially uncorrelated between pulses (its coherence time in slow time is much shorter than the integration window). Second, the cross-channel coherence: as derived in Section 1.6, an interferer entering the H receive antenna is identical in H_{HH} and H_{HV} , producing near-unity cross-channel coherence at the contaminated (f, τ) positions.

2.4. Effects on SAR Data and Science Products

The consequences of unmitigated RFI span from the most basic image quality metrics to the calibrated geophysical products that represent the mission’s scientific output.

In the focused SAR image the most visible effect of narrowband interference is a bright linear streak oriented along the azimuth direction, located at the range position corresponding to the apparent two-way delay of the interferer’s emission. The range position is determined by the frequency of the interferer relative to the centre of the SAR bandwidth: a tone at the edge of the receiver window maps to a range position offset from scene centre, while a tone at band centre appears superimposed on the nadir echo. This streak can mask legitimate scatterers at those range positions and, at sufficient RFI power, saturates the dynamic range of adjacent pixels through sidelobe leakage of the compressed interference peak. [8, 10]

Wideband interference produces a qualitatively different artefact: a spatially diffuse elevation of the noise floor across all range positions within the burst-contaminated azimuth

segment. This raises the effective NESZ locally, reducing the sensitivity of the backscatter measurement and making weak returns indistinguishable from the noise pedestal. For a TVWB interferer with $\text{SIR} = -15$ dB the local effective NESZ degrades by approximately 15 dB relative to the thermal noise floor. [8, 10]

The radiometric consequence is a bias in the estimated backscatter coefficient $\hat{\sigma}^0$. Because the interference power adds incoherently to the image intensity, it appears as a positive additive offset: $\hat{\sigma}^0 = \sigma_{\text{true}}^0 + \sigma_r^0$, where σ_r^0 is the interference-equivalent normalised power. This bias is spectrally and spatially non-uniform (it concentrates at the frequency and range positions where the interferer is most active) and therefore corrupts the radiometric calibration of all derived products including above-ground biomass maps that rest on calibrated σ^0 measurements.

Additionally, the most directly mission-critical effect is the degradation of interferometric and polarimetric coherence. During the interferometric Phase, an eventual RFI source that is statistically independent between two repeat-pass acquisitions – as is the case for terrestrial radio services that change state between Biomass overflights separated by days – contributes incoherently to the interferogram. The resulting coherence loss can be expressed through the multiplicative term:

$$|\gamma_{\text{RFI}}| = \frac{1}{1 + \text{SIR}^{-1}} \quad (2.4)$$

which decreases monotonically from unity (no interference) toward zero as the SIR worsens. For $\text{SIR} = 0$ dB, coherence is already halved; for $\text{SIR} = -10$ dB it drops to 0.09. Since the Cramér-Rao bound on interferometric phase noise (Equation (1.23)) grows rapidly as coherence decreases, even moderate RFI can render interferometric height estimates unusable. For the tomographic processing phase, where the joint processing of N passes is required and each coherence pair must exceed a minimum threshold, the loss of even a few passes to severe RFI contamination constrains the available baseline span and degrades the achievable elevation resolution. [10]

The polarimetric channel asymmetry is another practically important effect. Because the co-pol (HH, VV) channels carry stronger backscatter than the cross-pol (HV, VH) channels over most natural targets, an interferer of fixed power represents a higher effective SIR degradation for the cross-pol channels: the same interference power that produces $\text{SIR} = 0$ dB relative to HH may produce $\text{SIR} = -6$ to -10 dB relative to HV. Cross-pol measurements – which are essential for polarimetric decompositions and forest biomass retrieval at P-Band – are therefore systematically more vulnerable to RFI than co-pol measurements, a fact that single-channel detection methods fail to address because they

apply identical thresholds regardless of channel SNR. [9]

2.5. Detection and Suppression: State of the Art

The RFI problem in spaceborne SAR has received substantial attention since the deployment of L-Band systems such as JERS-1 and ALOS-PALSAR, which operate in a band similarly congested to P-Band. The literature subdivides the problem into detection (identifying which data samples are contaminated) and suppression (removing or attenuating the contamination), and most practical implementations address both sequentially using a detect-then-filter architecture identical in structure to the pipeline employed in this thesis.

The most widely adopted approach in the single-channel setting works in the (f, τ) domain and distinguishes TSNB and TVWB interference by the direction of their spectral persistence. Meyer et al. [8] formalised this into a quantitative framework: TSNB tones are detected by comparing the spectral power at each frequency bin against a Gaussian background estimated from the trimmed azimuth mean over a sliding window of ~ 256 lines, using a Fisher Z-test at a controlled false-alarm rate. TVWB events are detected by averaging the power over a window of ~ 100 range-frequency bins and testing for azimuth-direction power spikes above a detrended background. Both detection steps produce noisy binary masks that are then denoised using a Hough-line transform, which retains only connected linear features above a minimum length threshold and discards isolated pixels. The combined mask is applied as a spectral notch: contaminated samples are set to zero. This pipeline operates independently on each polarimetric channel and is the standard reference against which multi-channel methods are compared.

The primary limitation of single-channel spectral detection is its dependence on an accurate background power estimate. When the fraction of contaminated bins is large, the trimmed mean used as the background reference becomes biased, causing genuine signal bins to be flagged and removed. Empirically, when more than approximately 10–15% of the range-frequency bins are contaminated, the effective performance of this approach degrades noticeably: the cleaned channel may have a lower SNR than the corrupted input in the worst case, because the notch removes more valid signal than interference. A second limitation is the per-channel independence: the HV channel, with its lower scene SNR, is harder to detect in relative to a common interference level, so its mask will generally be less complete than the HH mask even for the same interferer.

Alternative single-channel strategies include parametric sinusoidal estimators (ESPRIT, MUSIC-type spectral estimators), which offer finer frequency resolution than the DFT-

based approach but at substantially higher computational cost, and adaptive notch filters implemented as IIR structures directly on the time-domain data. The latter avoid the Gibbs artefacts associated with spectral zeroing but require careful pole placement to avoid instability and do not offer the same sidelobe suppression as frequency-domain windowed notches. [7]

The recognition that polarimetric SAR offers a structural advantage for interference detection is more recent. Uozumi et al. [14] observed that the cross-channel coherence between co-pol and cross-pol channels, which is near zero over natural terrain, elevates to near unity wherever a shared-receive-antenna interference source is present, providing a detection signal that is entirely independent of the absolute power level. This observation forms the theoretical basis of the detection algorithm in Chapter 3: rather than comparing powers to a background model, it compares the cross-channel coherence to the null hypothesis of statistical independence.

For suppression, the Wiener-based interference cancellation framework (Widrow et al., 1975 [16]; Haykin [3]; Lord and Inggs [6]) provides the theoretically optimal linear canceller given a primary channel and a correlated reference. In the polarimetric SAR context, Wang et al. [15] applied a sidelobe cancellation filter using one polarimetric channel as the auxiliary reference to cancel interference in the other, showing that the coherent subtraction can achieve better suppression than notch filtering when the interference is coherent between channels and the weight estimate is accurate. The major limitation of their method (which the Wiener-validated method in this thesis addresses explicitly) is that the optimal weight varies with frequency and azimuth position, and a single global weight estimate degrades performance wherever the local coupling deviates from the average. The architecture described in Chapter 3 addresses this by combining local weight estimation with the safety constraint that cancellation is applied only where the detection mask has already identified a genuine RFI event.

Approaches based on subspace projection or low-rank interference modelling have been proposed for scenarios where the interference is known to occupy a small number of spectral modes. These methods offer strong suppression at low implementation cost when the interference model matches the data, but are sensitive to model mismatch — a problem in the variable P-Band environment where different interferers may be active on different acquisitions. The practical choice between model-based and data-driven approaches therefore depends on how reliably the interference structure can be characterised a priori; in the Biomass commissioning context, given the still limited knowledge and reliability of the measurements, and for the purpose of developing a method that does not rely on prior knowledge of the RFI distribution data-driven methods are preferred.

The literature on RFI in tomographic SAR is comparatively sparse. Tebaldini and Quegan [9, 11] noted that coherent interference appearing consistently across multiple passes produces a spurious peak in the tomographic reflectivity profile at an apparent elevation that cannot be distinguished from a real scatterer without cross-validation. The multi-pass covariance matrix $\mathbf{R}$ (Section 1.7) will contain a rank-1 interference contribution if the same interferer is active in all N passes at the same (f, τ) position, or a noise-like contribution (additional diagonal term) if the interferer is statistically independent across passes. Both modes of corruption affect the tomographic inversion, though in different ways: the rank-1 case creates a spurious scatterer, while the noise-like case raises the elevation-domain noise floor. The methods developed in this thesis, applied on a per-pass basis before multi-pass covariance estimation, would be able to address both modes: persistent interference would be detected and suppressed in each individual pass, and the cleaned covariance matrix could then be available for tomographic processing without the interference-induced bias.

3 | The Coherence-Based RFI Detector

The previous chapter established that radio frequency interference, once transformed to the range-frequency / azimuth-time domain (f, τ) , occupies compact and geometrically structured regions: horizontal stripes for Time-Stationary Narrowband (TSNB) sources and vertical stripes for Time-Varying Wideband (TVWB) sources. It also established, in Section 1.6, that an interferer entering a given receive chain contributes an identical waveform to both polarimetric channels served by that chain. This chapter combines the two observations into the detection algorithm that constitutes the first main contribution of this work: a detector that declares a (f, τ) cell as RFI-contaminated whenever the cross-channel coherence between the two channels sharing a receive port rises significantly above its clean-scene baseline. The approach follows the intuition first reported by Uozumi et al. [14], here extended with a complete mask-construction pipeline and adapted to the Biomass quad-polarimetric data format.

The chapter is organised as follows. Section 3.1 presents the physical motivation behind the effectiveness of the coherence statistic. Section 3.2 states and discusses the hypotheses under which the detector is expected to work, with particular attention to the timing relationship between the two channels of a receive pair. Section 3.3 derives the behaviour of the coherence estimator under the null and alternative hypotheses and explains how the detection thresholds were chosen. Section 3.4 describes the complete processing pipeline as implemented, Section 3.5 describes the tuning parameters for this method and their choice rationale, and Section 3.6 contrasts it structurally with the reference single-channel detector of Meyer et al. [8]. Section 3.7 closes the chapter with the known limitations and the ideal operating conditions of the method.

3.1. Physical Motivation

The detector rests on a simple physical asymmetry between the two components of the received data: the SAR backscatter is generated by the scene and is therefore shaped by

the scattering physics of the terrain, while the interference is generated by an external transmitter and is therefore shaped by the receiving hardware. These two origins leave opposite signatures in the cross-channel statistics of a receive pair.

Consider the H-receive pair, formed by the channels H_{HH} and H_{VH} (transmit polarisation first, receive polarisation second, as defined in the Introduction; the same reasoning applies verbatim to the V-receive pair H_{VV} , H_{HV}). Over natural distributed terrain with approximate azimuthal symmetry, the co-pol and cross-pol backscatter components are statistically independent: the mechanisms producing co-pol returns (surface and double-bounce scattering) are physically distinct from those producing cross-pol returns (volume scattering and multiple bounces with orientation diversity), and their ensemble cross-correlation vanishes, $E[S_{HH}S_{VH}^*] = 0$ (Section 1.6.2). The clean-scene cross-channel coherence of a receive pair is therefore near zero, with a residual magnitude set only by the finite size of the estimation window.

An external interferer behaves in the opposite way. The interfering wavefront arrives at the antenna with a fixed polarimetric state and is split by the receive hardware into the two channels served by the same physical port and both channels therefore record the *same* interference waveform. Writing the H-receive pair in the (f, τ) domain as in Equations (1.18)–(1.19):

$$H_{HH}(f, \tau) = S_{HH}(f, \tau) + r_H(f, \tau) + n_{HH}(f, \tau), \quad (3.1)$$

$$H_{VH}(f, \tau) = S_{VH}(f, \tau) + r_H(f, \tau) + n_{VH}(f, \tau), \quad (3.2)$$

the cross-correlation of the pair is dominated by the common term wherever the interference power is significant, and the cross-channel coherence (Equation (1.14)) transitions from near zero toward unity.

Two properties make this transition an exceptionally convenient detection signal. First, it is *waveform-agnostic*: the derivation above made no assumption on the interferer's modulation, bandwidth, frequency, or duty cycle. Any external signal that couples into both channels of a receive pair raises the coherence, regardless of whether it would be classified as TSNB, TVWB, or something in between. Second, it is *power-normalised*: the coherence is a ratio, bounded in $[0, 1]$, whose clean-scene baseline does not depend on the local backscatter level. Unlike amplitude-based detectors, no background power model must be estimated from the (possibly already contaminated) data, and the same threshold applies identically over bright urban areas and dark water surfaces. This is the structural reason why, as will be quantified in Chapter 4, the coherence detector retains its sensitivity on the weak cross-pol channels where amplitude statistics become unreliable.

3.2. Working Hypotheses

For the qualitative argument of the previous section to translate into a reliable decision statistic, four hypotheses must hold. They are stated here explicitly, because each of them maps onto a specific design choice of the pipeline of Section 3.4 and onto a specific failure mode discussed in Section 3.7.

3.2.1. H1 – Scene independence within the receive pair

The backscatter components S_{HH} and S_{VH} must be statistically independent, so that the clean-data coherence baseline is governed only by estimation noise. As discussed in Section 1.6.2, this follows from reflection symmetry of natural distributed targets and holds well over forests – the primary Biomass science target – and most natural terrain. The hypothesis weakens over scenes with strong oriented structures (urban areas, agricultural rows, slopes oriented along specific directions), where a genuine physical co-pol/cross-pol correlation exists. Importantly, the detector operates on *defocused* data, where the contribution of each scatterer is smeared along the full synthetic aperture and across range by the inverted chirp: even where a residual scene correlation exists in the focused image, the defocusing dilutes it over a large (f, τ) support, lowering the baseline coherence and protecting the null hypothesis.

3.2.2. H2 – Common coupling of the interference

The interference recorded by the two channels of the pair must be the same waveform up to a complex scalar, as in Equations (3.1) and (3.2). This is guaranteed by the receive architecture whenever the two channels share the same physical antenna port and the front-end transfer functions are stable over the estimation window. The hypothesis would fail for an interferer whose polarisation state changes rapidly compared to the window duration – a physically uncommon situation for the fixed terrestrial transmitters described in Chapter 2 – or for internally generated (channel-specific) disturbances, which by construction this detector cannot see.

3.2.3. H3 — Negligible inter-channel acquisition delay

In a quad-polarimetric SAR the transmit polarisation alternates between H and V on successive pulses, while both receive chains record every echo. The two channels of a receive pair are therefore not acquired simultaneously: the H-receive samples of H_{HH} come from the H-transmit pulses and those of H_{VH} from the V-transmit pulses, so corresponding

azimuth samples of the pair are separated by one pulse repetition interval,

$$\Delta\tau = \frac{1}{\text{PRF}}, \quad (3.3)$$

while each individual channel is sampled at $\text{PRF}/2$. The scene component is insensitive to this offset since it is uncorrelated between the channels regardless (according to hypothesis H1). The interference, however, is only *approximately* common: channel one records $r_{\text{H}}(f, \tau)$ and channel two records $a r_{\text{H}}(f, \tau + \Delta\tau)$.

The detector therefore requires that the interference, as observed in a single range-frequency bin, decorrelates slowly compared to $\Delta\tau$:

$$\Delta\tau \ll \tau_c^{\text{RFI}}, \quad (3.4)$$

where τ_c^{RFI} is the coherence time of the bin-filtered interference envelope. Under this condition the lagged samples are nearly identical, $r_{\text{H}}(f, \tau + \Delta\tau) \approx r_{\text{H}}(f, \tau)$, and the common-signal argument of Section 3.1 applies in full.

It is worth examining the consequences when Equation (3.4) does not hold to a satisfactory level. As the decorrelation time approaches the PRI, this would produce a residual misalignment between the complex values of RFI affected pixels, which bounds the maximum coherence the detector can observe and caps the cancellation depth achievable by the coherent removal method of Chapter 5, motivating the azimuth-alignment improvement proposed in Section 5.6, where the known acquisition timing is used to re-align the two channels before the coherence is computed.

3.2.4. H4 – Local interference dominance

The coherence rises toward unity only where the interference power is a significant fraction of the total power in *both* channels of the pair (a statement made quantitative in the next section). The detector is therefore a detector of *contaminated cells*, not of arbitrarily weak interference: a source buried far below the scene return in a given bin will not be flagged. This behaviour is in fact desirable, since cells where the interference is negligible do not need correction, and removing them would only destroy useful signal. The relevant operational consequence is that detection sensitivity is naturally higher on the cross-pol channels, whose lower backscatter makes a given interference power a larger fraction of the total (Section 2.4).

3.3. The Detection Statistic

3.3.1. Coherence under contamination

The behaviour of the population coherence under the signal model [(3.1), (3.2)] admits a compact closed form. Assuming the received signal to be stationary within an estimation window, let $P_{S1} = E[|S_{HH}|^2]$, $P_{S2} = E[|S_{VH}|^2]$, $P_r = E[|r_H|^2]$, and P_{n1}, P_{n2} the noise powers. Since scene, interference, and noise are mutually uncorrelated and the scene components are uncorrelated with each other (H1), the cross-moment reduces to the interference term, $E[H_{HH}H_{VH}^*] = P_r$, while the channel powers are $P_1 = P_{S1} + P_r + P_{n1}$ and $P_2 = P_{S2} + P_r + P_{n2}$. Substituting into the definition of Equation (1.14),

$$|\gamma| = \frac{P_r}{\sqrt{P_1 P_2}} = \sqrt{\rho_1 \rho_2}, \quad \rho_1 = \frac{P_r}{P_1}, \quad \rho_2 = \frac{P_r}{P_2}, \quad (3.5)$$

where $\rho_1, \rho_2 \in [0, 1]$ are the fractions of total power contributed by the interference in the two channels. Equation (3.5) is the quantitative version of hypothesis H4: the coherence is the geometric mean of the per-channel interference fractions. It approaches unity when the interference dominates both channels, equals 0.5 when the interference carries one quarter of the power in each channel (i.e. a per-channel SIR of about +5 dB), and decays to zero as the contamination becomes negligible. Notably, expression (3.5) does not contain the absolute power level: scaling the scene, the interference, and the noise together leaves the statistic unchanged, which was expected after the explanations of Section 3.1.

3.3.2. The estimator and its null distribution

The population coherence is estimated, as in Equation (1.15), from a sliding window of W samples in the (f, τ) plane:

$$\hat{\gamma}(f, \tau) = \frac{\left| \sum_{i \in \mathcal{W}(f, \tau)} H_{\text{co}}(i) H_{\text{x}}^*(i) \right|}{\sqrt{\sum_{i \in \mathcal{W}(f, \tau)} |H_{\text{co}}(i)|^2 \sum_{i \in \mathcal{W}(f, \tau)} |H_{\text{x}}(i)|^2}}, \quad (3.6)$$

where $\mathcal{W}(f, \tau)$ is a window of W contiguous cells centred on (f, τ) and the subscripts “co”/“x” denote the two channels of the active receive pair (respectively co-polarized and cross-polarized). The implementation evaluates the three sums by 2-D convolution with a constant kernel, so the full coherence map is obtained at a cost linear in the number of pixels.

Under the null hypothesis (clean data, $\gamma \approx 0$) the estimator does not return zero: with

W independent complex samples, $\hat{\gamma}$ has a known bias with expected magnitude decaying as $E[\hat{\gamma} | \gamma = 0] \simeq \sqrt{\pi}/(2\sqrt{W})$, and the probability that the estimate exceeds a threshold T follows

$$\Pr(\hat{\gamma} > T | \gamma = 0) = (1 - T^2)^{W-1}, \quad (3.7)$$

a classical result for the sample coherence of independent circular Gaussian processes [13]. Equation (3.7) is the design tool of the detector: it links the window size W and the threshold T to the per-pixel false-alarm probability before any post-processing is done. With the baseline choice $W = 10$ samples, suggested by the original formulation [14] and confirmed by the simulations of Chapter 4, a threshold of $T = 0.5$ yields a raw per-pixel false-alarm rate of $(1 - 0.25)^{15} \approx 1.3\%$, subsequently reduced by the line-detection and rejection stages described below. Larger windows tighten the null distribution and allow lower thresholds (better sensitivity through H4), at the cost of spatial-spectral resolution of the mask and of requiring interference stationarity over a longer support.




Figure 3.1: Double estimation window working principle

3.3.3. Two oriented estimators

A single isotropic window would dilute the elongated RFI signatures established in Section 2.3 into their clean surroundings. The detector therefore computes Equation (3.6) twice, with two mutually orthogonal windows that match the expected morphologies:

- $\hat{\gamma}_{\text{TSNB}}$ uses a $1 \times W$ window extended along the *azimuth* direction. A persistent narrowband tone occupies the same range-frequency bin over many consecutive pulses,

so integrating along azimuth accumulates the common component coherently while averaging the independent scene contributions toward zero;

- $\hat{\gamma}_{\text{TVWB}}$ uses a $W \times 1$ window extended along the *range-frequency* direction. A wideband burst occupies many frequency bins on the same (or few) azimuth lines, so the integration runs across frequency.

Each estimator is maximally sensitive to its target class and insensitive to the other, mirroring – in the coherence domain – the two-detector architecture of the Meyer method, but replacing its two background-power models with a single physical statistic.

3.4. The Processing Pipeline

This section walks through the implemented pipeline in execution order, while Algorithm 3.1 provides the equivalent pseudo-code.

3.4.1. Defocusing and domain transform

The detector operates on raw data, where the RFI structure is maximally compact (Section 2.3). Since the available input is the focused Single Look Complex (SLC) product, the first step inverts the focusing chain: an inverse Range-Doppler processor removes the azimuth matched filter and deconvolves the range chirp through an inverse matched filter, returning the raw-like matrix $h(t, \tau)$ for each channel. A 1-D FFT along fast time then yields the working representation $H(f, \tau)$. The fidelity of this defocus–refocus round trip, and its role in the simulation environment, are characterised in Section 4.2.

3.4.2. Pair-wise coherence and thresholding

The oriented coherence maps $\hat{\gamma}_{\text{TSNB}}$ and $\hat{\gamma}_{\text{TVWB}}$ are computed for the H-receive pair $(H_{\text{HH}}, H_{\text{VH}})$ and V-receive pair $(H_{\text{VV}}, H_{\text{HV}})$. Each map is compared against its threshold to produce raw binary maps. The exact threshold values used in this work, together with a sensitivity analysis, are reported in Chapter 4.

A *band restriction* is optionally applied at this point: cells whose range frequency falls outside the transmitted system bandwidth B (oversampled products carry guard bands at $|f| > B/2$) cannot contain scene signal by construction, so their statistics carry no scene–interference contrast. They are excluded from the mask logic and handled separately, since interference falling there is removed at no cost.

3.4.3. Line detection

The raw binary maps are noisy: isolated false alarms from the tail of (3.7) coexist with small gaps inside genuine RFI stripes. The structural reasoning that RFI forms straight, axis-aligned features in the (f, τ) plane is enforced by a line-detection stage that retains only connected horizontal (TSNB map) or vertical (TVWB map) runs longer than a minimum length $L_{\min}$, discarding everything else. Because a genuine tone persists over hundreds of azimuth lines while a genuine burst spans a sizeable fraction of the band, $L_{\min}$ can be set generously (with a proportionally larger value for the azimuth-elongated TSNB features) at a negligible cost in missed detections, while suppressing the isolated false alarms almost entirely. This stage plays the same role as the Hough-transform denoising in the Meyer pipeline, at a lower computational cost since the orientation of the sought lines is known a priori.

3.4.4. Mask application and refocusing

For each processed pair, the mask derived from the TSNB detection and TVWB detection are combined with an OR logic, since they are optimized to detect independent phenomena. Each channel couple sharing the same receive polarity will also share the same detection mask, and the removal stage can be applied: in the baseline configuration of this chapter, removal is performed by notching, meaning that the flagged cells are zeroed:

$$H^{\text{notch}}(f, \tau) = H(f, \tau) (1 - M(f, \tau)), \quad (3.8)$$

after which the data are inverse-transformed and refocused. The notch removes the interference exactly but also annihilates the scene content of the flagged cells, an information loss whose consequences, and whose mitigation through coherent cancellation, are the subject of Chapter 5.

Algorithm 3.1 Coherence-based RFI detection and notching

- 1: Defocus the four focused channels: $I_{pq} \rightarrow h_{pq}(t, \tau)$, $pq \in \{\text{HH}, \text{HV}, \text{VH}, \text{VV}\}$
 - 2: Range FFT: $H_{pq}(f, \tau) = \mathcal{F}_t\{h_{pq}\}$
 - 3: **for** each receive pair $(\text{co}, \text{x}) \in \{(\text{HH}, \text{VH}), (\text{VV}, \text{HV})\}$ **do**
 - 4: Compute $\hat{\gamma}_{\text{TSNB}}$ with a $1 \times W$ azimuth window (Equation (3.6))
 - 5: Compute $\hat{\gamma}_{\text{TVWB}}$ with a $W \times 1$ range-frequency window
 - 6: Threshold: $B_{\text{TSNB}} = \hat{\gamma}_{\text{TSNB}} > T + 0.25$; $B_{\text{TVWB}} = \hat{\gamma}_{\text{TVWB}} > T$
 - 7: Apply band restriction to $B_{\text{TSNB}}, B_{\text{TVWB}}$
 - 8: Line detection: keep horizontal runs $\geq 15L_{\min}$ in B_{TSNB} , vertical runs $\geq L_{\min}$ in B_{TVWB}
 - 9: OR-combine the two detection masks into the candidate mask
 - 10: Apply the shared mask M : notch (Equation (3.8)) on both channels currently under processing
 - 11: **end for**
 - 12: Inverse range FFT and refocus each channel
-

3.5. Tuning Parameters

The algorithm has five primary parameters that must be set by the user. Their physical meaning and typical operational ranges are described below.

3.5.1. Coherence Threshold γ_{th}

The coherence threshold $\gamma_{\text{th}} \in (0, 1)$ is the primary sensitivity control. A lower threshold increases detection sensitivity (reduced missed-detection probability P_{MD}) at the cost of higher false-alarm probability P_{FA} , and vice versa. The operating point is selected based on the distribution of $|\hat{\gamma}|$ under the null hypothesis $\mathcal{H}_0$, under which the signals received by the co- and cross-polarized channels are uncorrelated.

Assuming $\mathcal{H}_0$ to be true, the complex coherence between two uncorrelated channels estimated over L samples follows a known distribution whose variance decreases as $1/L$ [13]. For large L , the magnitude $|\hat{\gamma}|$ is approximately distributed as:

$$\mathbb{E}[|\hat{\gamma}|^2]_{\mathcal{H}_0} \approx \frac{1}{L} \quad (3.9)$$

suggesting that the 3σ false-alarm limit scales as $\gamma_{\text{th}} \approx 3/\sqrt{L}$. In practice, γ_{th} is tuned empirically on a set of clean reference acquisitions. Typical operational values lie in the

range $[0.4, 0.7]$ depending on the window sizes and the scene characteristics.

3.5.2. TSNB Window Length W_F

W_F is the number of azimuth lines integrated in the TSNB estimator. A larger W_F reduces the estimation variance of $\hat{\gamma}_{\text{TSNB}}$ (more independent samples) but reduces sensitivity to short-duration narrowband bursts. For the Biomass mission geometry (PRF ≈ 1.5 kHz), W_F values in the range 8–256 lines were considered, corresponding to integration times of 0.02–0.17 s.

3.5.3. TVWB Window Width W_T

W_T is the number of frequency bins integrated in the TVWB estimator. A larger W_T provides better averaging against frequency-domain noise but reduces spectral resolution (adjacent narrowband features within the same window may be merged). The values assessed during tuning span 8–64 bins, corresponding to W_T/f_s intervals of a few MHz.

3.5.4. Minimum Line Length $\ell_{\min}$

$\ell_{\min}$ controls the aggressiveness of the line-length filter: isolated false-alarm pixels are suppressed, but a longer required length implies that short RFI bursts (duration less than $\ell_{\min}$ samples) will not be detected. A practical minimum is $\ell_{\min} \geq W_F$ for TSNB and $\ell_{\min} \geq W_T$ for TVWB, ensuring that any detected feature spans at least the integration window duration.

3.6. Comparison with the Meyer Detector

The detector of Meyer et al. [8], summarised in Section 2.5, is the established single-channel reference against which the proposed method is benchmarked throughout this work, and a structural comparison clarifies what is being traded. Both methods operate in the same (f, τ) domain, both target the same two interference classes with two dedicated sub-detectors, and both regularise their raw decisions with a line-based morphological stage before applying a notch. The difference is concentrated in the decision statistic itself. The Meyer detector tests *amplitude* against a background model: the TSNB branch compares each frequency bin of the azimuth-averaged spectrum (over $N_{\text{az}} = 256$ lines) against a Gaussian background whose mean and deviation are estimated by a trimmed mean, through a one-tailed Fisher Z-test at a controlled false-alarm level; the TVWB branch applies an analogous test to the detrended azimuth power profile averaged over

$N_{\text{rng}} = 100$ frequency bins. The coherence detector tests *cross-channel correlation* against the null hypothesis of independence, requiring no background estimate at all.

Three consequences follow. First, the Meyer statistic inherits the vulnerability of every background-based test: when the contaminated fraction of bins grows large, the trimmed mean is biased upward, the test loses calibration, and clean bins begin to be flagged, or contaminated pixels stay below the flagging threshold – the empirical breakdown observed beyond roughly 10–15% contamination (Section 2.5). The coherence statistic has no background to corrupt: heavy contamination simply enlarges the region where Equation (3.5) is close to one. Second, the Meyer detector processes each polarimetric channel in isolation and therefore demands that each channel individually carry enough contrast for its Z-test, which penalises the low-backscatter cross-pol channels; the coherence detector lets the stronger channel of each receive pair lend statistical support to the weaker one, and its shared mask guarantees polarimetric consistency of the correction. On the other hand, the per-channel independence of the Meyer masks is also a form of robustness that the shared mask gives up: a false alarm of the coherence detector contaminates both channels at once. The false-alarm rejection stage based on line detection of Section 3.4 exists precisely to keep this failure mode rare. On the computational side, the coherence pipeline reduces to separable convolutions and connected-component analysis, scaling linearly with the number of pixels, whereas the Hough accumulation of the reference method is the dominant cost in its pipeline.

	Meyer et al.	Coherence Detector
Decision statistic	Single-channel power vs. Gaussian background (Fisher Z-test)	Cross-channel coherence vs. empirical threshold
Channel Processing	Each processed independently	Receive pairs processed jointly; all four share one mask
Background signal model	Trimmed-mean spectral estimate	None required
False-alarm control	Z-test level + Line Detection	Threshold from Eq. (3.7) + Line Detection
Limitations	Background bias at high contamination	Higher False alarm rate from shared mask; Stricter operating hypothesis

Table 3.1: Structural comparison between the reference single-channel detector and the proposed coherence-based detector.

3.7. Limitations and Ideal Operating Conditions

The hypotheses of Section 3.2 delimit the envelope inside which the detector is expected to perform well, and each of them identifies a corresponding limitation.

The method is inherently polarimetric: it requires at least one receive-sharing channel pair, and is therefore inapplicable to single-polarisation acquisitions, where amplitude-based methods such as Meyer’s remain the only option. Within a pair, sensitivity is governed by Equation (3.5): interference well below the scene level in either channel of the pair is not flagged, which is acceptable for the cells that need no correction but implies that the detector should not be read as an exhaustive interference census. Scenes that violate azimuthal symmetry raise the coherence baseline and erode the detection margin; additionally, the defocused-domain operation and the morphological constraints mitigate, but cannot eliminate, this effect over strongly oriented man-made scenes. The timing hypothesis H3 caps the observable coherence for rapidly modulated interferers, with the consequences and the remedies discussed in Section 5.6. Finally, the shared mask makes false alarms more costly than in per-channel schemes, a deliberate trade accepted in exchange for the polarimetric consistency of the correction and contained by the rejection pipeline.

The ideal operating conditions follow by complement: quad-polarimetric (or dual-pol with shared receive) acquisitions over natural distributed scenes; interference that couples into the receive chains as an external wavefront, with per-cell power at least comparable to the scene return; interferers whose bin-filtered envelopes are stable at the pulse-to-pulse scale; and stationarity of the interference over the coherence estimation window. In order to verify whether the biomass mission environment is optimal for this detection method, the real-data campaigns of Chapter 5 are necessary to validate the simulated performances presented in the following chapter.

4 | Simulation Environment and Detector Validation

A detector can only be validated where the truth is known. On real interference-affected acquisitions there is no access to the clean scene, no exact knowledge of which cells are contaminated, and no control over the interference parameters: any performance figure computed there is, at best, indirect. Before applying the coherence-based detector of Chapter 3 to real Biomass data, it was therefore validated in a fully controlled simulation environment, in which synthetic interference with known parameters is injected into a clean SAR image and the output of every method can be compared against an exact ground truth. The simulation serves three purposes: it verifies that the detector behaves according to the theory of Section 3.3; it allows a fair, quantitative comparison with the reference Meyer detector under identical conditions, with both methods followed by the same removal step (notching) so that only the detection stage is being compared; and it provides the tuning of the pipeline parameters (window size, thresholds, minimum line lengths) that are then frozen before the real-data experiments of Chapter 5.

4.1. Architecture of the Simulation

The simulation is built around a clean, RFI-free focused quad-polarimetric SLC tile, used as the reference scene. The reference scene is a clean, RFI-free focused quad-polarimetric SLC tile extracted from a Biomass commissioning-phase acquisition and cleaned of invalid samples before use. That it is effectively interference-free is confirmed a posteriori by its cross-pol image coherence, whose mean over the HH–VH pair is 0.091: such value is expected for independent scene backscatter over natural terrain (Section 3.1) and far below the levels reached under contamination. The tile spans the full system bandwidth ($f_s = 7.565$ MHz, slightly over-sampled with respect to the system 6 MHz bandwidth B) and the processed image azimuth size is large enough to apply the azimuth-pattern modes of the injector while keeping computational cost low enough for fast execution and eventual parameter sweeps for tuning purpose. Denoting the clean focused channels

by I_{pq}^{clean} , the experiment proceeds along the chain illustrated in Figure 4.1: the clean channels are defocused to the raw-equivalent domain, synthetic interference $r(t, \tau)$ with prescribed type and power is added there and the corrupted data are refocused into the corrupted focused images that constitute the input of every method under test. Each method returns cleaned SLC images, which are then evaluated against I_{pq}^{clean} and against the known injected interference r .




Figure 4.1: Architecture of the controlled validation environment. The injection takes place in the raw-image domain, where physical interference actually enters the receiver

4.1.1. The defocus–refocus operator pair

The reliability of the simulations heavily depends on the invertibility of the transformation between the focused product and the raw-equivalent domain. The defocusing operator implements the exact inverse of a Range-Doppler focusing chain: an azimuth FFT brings the image to the range-Doppler domain; the azimuth matched-filter phase is removed per range gate; the range cell migration is re-introduced by sinc interpolation along range; and the range chirp is de-convolved in the frequency domain. The chirp deconvolution deserves a note: a plain inverse filter $1/G(f)$ would amplify without bound the frequencies where the chirp spectrum $G(f)$ rolls off, so a Tikhonov-regularised inverse is used instead,

$$G_{\text{reg}}^{-1}(f) = \frac{G^*(f)}{|G(f)|^2 + \epsilon \max_f |G(f)|^2}, \quad (4.1)$$

with a small regularisation fraction ϵ (of the order of 10^{-3}), which behaves as the exact inverse inside the transmitted band and gracefully limits the gain outside it. The refocusing operator applies the same steps in reverse order, reusing the geometric and timing

information stored at defocus time so that the round trip is sample-aligned.

The round-trip fidelity was verified by defocusing and immediately refocusing the clean tile without any injection: the reconstruction error is dominated by the regularised band edges and by the interpolation of the range migration correction.




Figure 4.2: repeated defocus-refocus effect and introduced distortions

This residual fixes the noise floor of the simulation: no mitigation method can be measured beyond the fidelity of the round trip itself, and all metrics reported in this chapter are to be read against this floor. Measured on the H-receive channels and then averaged, this round-trip residual amounts to $\text{NMSE} = 0.196$ (below which the NMSE of any method cannot fall). This value is the reference for the reconstruction metrics of Table 4.2: taking as example the co-pol channel, its results are more distinct than it appears just from the absolute values. The corrupted input sits at $\text{NMSE} = 0.381$, which is 0.185 above the floor; of this recoverable error the coherence detector removes roughly 90%, reaching 0.215 (only 0.019 above the floor) while the single-channel reference removes under a third and remains at 0.329. The coherence reconstruction of the clean properties is therefore floor-limited rather than interference-limited, coming extremely close to the simulation optimal result.

4.1.2. Synthetic interference injection

The injected interference follows the two analytical models of Section 2.2, implemented as parametric generators operating in the raw-equivalent domain.

TSNB generator. Each narrowband event is a complex tone of programmable base-band frequency ν_0 (with $|\nu_0| < f_s/2$), amplitude A_0 , and range support: the tone can occupy the full receive window or any sub-interval of it, with a tapered rectangular win-

down, reproducing the behaviour of transmitters discussed in Section 2.2. Along azimuth, the activity pattern $L_{az}(\tau)$ implements the cases of Equation (2.1): continuous activity (the physically standard case for a stationary emitter), periodic activity every N lines, or random activity with a programmable per-line probability; the amplitude can additionally be modulated along azimuth (constant, linear ramp, or sinusoidal) to emulate antenna-pattern weighting of the interferer as the platform passes by. Multiple tones with independent parameters can be superimposed, producing the multi-carrier environments typical of congested bands.

TVWB generator. Each wideband event is a burst confined to a programmable fast-time interval (start position and duration as fractions of the receive window) with an amplitude taper across the burst (rectangular, Hann, Hamming, or Blackman). The burst waveform implements the three variants of Equations (2.2)–(2.3): band-limited complex Gaussian noise of programmable bandwidth and centre frequency; a deterministic LFM chirp whose bandwidth is set by the burst duration, modelling coherent pulsed radars; or a periodic train of LFM sub-pulses with programmable sub-pulse width and repetition rate, modelling fast-scanning systems whose inter-pulse period is shorter than the SAR receive window. The azimuth pattern offers the same modes as the TSNB generator, with one addition specific to stochastic bursts: the noise realisation can be regenerated independently for every azimuth line or shared across lines (fully coherent burst).




Figure 4.3: example of the expected behaviour of the artificial RFI injection

Power calibration and channel coupling. The interference amplitude is set through the target Signal-to-Interference Ratio: given the measured mean scene power P_{sig} of the tile and the desired SIR_{dB} , the generator scales the event to $A_0 = \sqrt{P_{sig} / 10^{SIR_{dB}/10}}$, and the actually realised SIR is recomputed a posteriori from the injected samples and logged

together with the exact set of contaminated cells. The latter, converted to the (f, τ) grid, constitutes the ground-truth detection mask M^{GT} . Consistently with the receive-chain physics of Section 3.1, each interference realisation is injected identically into the two channels of a receive pair – the same waveform enters h_{HH} and h_{VH} , and an independent realisation enters h_{VV} and h_{HV} – so that hypothesis H2 of Section 3.2 is satisfied by construction and the experiments probe the remaining hypotheses and the pipeline itself.

4.2. Test Scenarios

The validation campaign sweeps the two interference classes separately and jointly across a grid of contamination severities. The structure of the campaign is as follows:

- **TSNB-only scenarios:** two continuous narrowband tones at 1.5 and 2.5 MHz baseband, each spanning the full range window with an always-on azimuth pattern, with both a High-SIR and Low-SIR case [$\text{SIR}_{\text{TSNB}} = 0$ dB and $\text{SIR}_{\text{TSNB}} = +5$ dB];
- **TVWB-only scenarios:** one band-limited complex-Gaussian-noise burst of 3 MHz bandwidth centred at -0.5 MHz, only injected for the second half of the fast time values and with random azimuth activity ($\sim 10\%$ of lines), and both High- and low-SIR case were investigated [$\text{SIR}_{\text{TVWB}} = -7$ dB and $\text{SIR}_{\text{TVWB}} = -2$ dB];
- **Mixed scenarios:** simultaneous TSNB and TVWB injection at $\text{SIR}_{\text{TSNB}} = +5$ dB and $\text{SIR}_{\text{TVWB}} = -2$ dB, reproducing the congested-band conditions expected at P-Band. This is the scenario analysed in detail in Section 4.4.

Every scenario is processed by the two pipelines under comparison – (i) Meyer detection followed by notching and (ii) coherence detection followed by notching – with the removal step deliberately kept identical, so that any performance difference is attributable to the detection stage alone.






Figure 4.4: Simulated RFI for the mixed case injected over a blank image background



(a) Isotropic coherence maps: Clean (upper-left), Corrupted with artificial RFI (upper-right), After Meyer algorithm (lower-left), After coherence algorithm (lower-right)



(b) Detection masks: comparison between the H-receive shared mask and HH-derived mask with Meyer's method

Figure 4.5: RFI Injection and detection comparison: coherence elevates toward unity exactly over the injected structures, as predicted by Equation (3.5).

4.3. Performance Metrics

Because the clean tile I^{clean} and the exact injected interference r are both known, the simulation supports three complementary families of metrics: a first category named *detection* metrics, which score the mask against the known contamination geometry; secondly, *per-*

channel reconstruction metrics, which score each cleaned image against its clean reference; and lastly *polarimetric* metrics, which score the recovered two-channel statistics that feed the interferometric and tomographic products of Section 1.7. All quantities are read comparatively across methods rather than as absolute figures of merit. Throughout, I^{out} is a method's focused output, I^{clean} the clean reference, $I^{\text{in}} = I^{\text{clean}} + r$ the corrupted input, and $\langle \cdot \rangle$ a spatial average.

4.3.1. Detection metrics

The detector returns a binary mask M in the range-frequency/azimuth-time (f, τ) plane. Because the injection is known, the *ground-truth mask* M^{gt} is defined as the set of cells in which the injected RFI power exceeds the clean signal power, i.e. the cells where the local signal-to-interference ratio is below 0 dB,

$$M^{\text{gt}} = \{(f, \tau) : |R(f, \tau)|^2 > |S(f, \tau)|^2\}, \quad (4.2)$$

with R and S the range-Fourier transforms of the injected RFI and of the clean signal. M^{gt} is built in each detector's native plane — the defocused (f, τ) plane for the coherence detector, the focused range-frequency plane for the single-channel detector — from the same injection, so each mask is scored on the grid it actually produces.

Detection probability P_d (true positive) and False Alarm (false-negative) rate P_{fa} . With $|\cdot|$ the cell count,

$$P_d = \frac{|M \cap M^{\text{gt}}|}{|M^{\text{gt}}|}, \quad P_{fa} = \frac{|M \setminus M^{\text{gt}}|}{|M^{\text{gt}}|}, \quad (4.3)$$

that is, the fraction of truly-contaminated cells that were flagged (P_d , higher is better) and the fraction of clean cells wrongly flagged (P_{fa} , lower is better).

Mask coverage. The overall flagged fraction $|M|/N$ (N the number of cells) is also tracked: under notch removal every flagged cell is a cell whose scene content is discarded, so over-masking translates directly into scene-information loss even when each flagged cell is genuinely (if negligibly) contaminated.

4.3.2. Per-channel reconstruction metrics

These quantify how faithfully each method returns a single corrupted channel to its clean reference; they are computed per channel on the focused output.

Normalised mean-square error (NMSE). The complex (amplitude-and-phase) error power, normalised by the clean power,

$$\text{NMSE} = \frac{\|I^{\text{out}} - I^{\text{clean}}\|_{\Omega}^2}{\|I^{\text{clean}}\|_{\Omega}^2}, \quad (4.4)$$

with 0 the ideal. Evaluating it over the whole image ($\Omega = \text{all}$), over the flagged cells ($\Omega = \text{flag}$) and over the complement ($\Omega = \text{unflag}$) separates damage *inside* the mask from damage leaked *outside* it.

SNR improvement (ΔSNR). The same quantity in decibels, oriented so that larger is better,

$$\Delta\text{SNR} = 10 \log_{10} \frac{\|I^{\text{clean}}\|^2}{\|I^{\text{out}} - I^{\text{clean}}\|^2} = -10 \log_{10} \text{NMSE}. \quad (4.5)$$

A method whose output figure falls below the input figure (equivalently $\Delta\text{SNR} < 0$, as for the heavily-corrupted cross-pol input) has introduced more distortion than it removed.

Amplitude RMSE. The magnitude-only error, in decibels relative to the clean power,

$$(\text{Amplitude RMSE})_{\text{dB}} = 10 \log_{10} \frac{\langle (|I^{\text{out}}| - |I^{\text{clean}}|)^2 \rangle}{\langle |I^{\text{clean}}|^2 \rangle}, \quad (4.6)$$

lower being better. Unlike NMSE it ignores phase, isolating the radiometric (σ^0) fidelity that drives backscatter-based biomass retrieval.

Reconstruction coherence (γ_{rec}). The interferometric coherence between I^{out} and I^{clean} (Equation (1.15)), spatially averaged over a square window:

$$\gamma_{\text{rec}} = \langle |\gamma(I^{\text{out}}, I^{\text{clean}})| \rangle \in [0, 1], \quad (4.7)$$

1 being a perfect reconstruction. Being a coherence, it is the fidelity notion most directly relevant to data feeding interferometric and tomographic processing (Section 1.7).

RFI-residual power. Since $I^{\text{in}} - I^{\text{clean}} = r$ is exactly the injected interference, the fraction of its power surviving in the output is

$$\text{RFI-residual} [\%] = 100 \cdot \frac{\|I^{\text{out}} - I^{\text{clean}}\|^2}{\|I^{\text{in}} - I^{\text{clean}}\|^2} = 100 \cdot \frac{\text{NMSE}}{\text{NMSE}_{\text{in}}}, \quad (4.8)$$

so 100% means no removal, 0% perfect removal, and values above 100% flag a method that distorted the data beyond the original contamination.

4.3.3. Polarimetric pair-recovery metrics

The cross-channel statistics of a receive pair $(a, b) = (\text{HH}, \text{VH})$ are the quantities that ultimately matter for polarimetric and Pol-InSAR products; they are recovered far less trivially than a single channel and are scored against the clean pair. Let γ_{ab} denote the cross-channel coherence (Equation (1.15)) and $C_{ab} = \langle I_a I_b^* \rangle$ the off-diagonal element of the polarimetric covariance (Section 1.6).

Mean image coherence. The spatial mean $\langle |\gamma_{ab}| \rangle$ of the cross-channel coherence on the *focused* cleaned images. The target is the clean value (≈ 0.095 for the natural scene used as background in the simulation); both over- and under-shoot are possible, so the relevant quantity is the distance to the clean figure.

Coherence-map RMSE. The per-pixel RMS difference between the cleaned and the clean coherence maps, $\langle (|\gamma_{ab}^{\text{out}}| - |\gamma_{ab}^{\text{clean}}|)^2 \rangle^{1/2}$, lower being better: it highlights the local coherence errors that the spatial mean averages away.

Unfocused coherence (overall mean and average at flagged cells). The same cross-channel coherence computed in the unfocused (f, τ) domain – in which the detector operates – with window $W = 10$ reported both as a whole-plane mean and averaged over the RFI flagged cells only. The latter is the most direct measure of how much common interference survives at the contaminated locations; the target is again the clean value.

Cross-covariance relative error. The relative error of the off-diagonal covariance element:

$$\varepsilon_C = \frac{|C_{ab}^{\text{out}} - C_{ab}^{\text{clean}}|}{|C_{ab}^{\text{clean}}|} \quad (4.9)$$

where 0 is the optimal value. This metric is the single number closest to what a polarimetric decomposition actually evaluates.

Runtime. The wall-clock time per scene, reported for completeness as the efficiency cost of each pipeline.

4.4. Results

4.4.1. Detection performance

Channel	Metric	Corrupted	Meyer	Coherence
HH	Mean γ	0.807	0.925	0.942
	RFI resid.	100%	19.4%	10.5%
VH	Mean γ	0.558	0.805	0.906
	RFI resid.	100%	17.2%	6.7%
TVWB detection	Precision $\left[\frac{N_{\text{correct detection}}}{N_{\text{tot detection}}} \right]$	-	0.978	0.796
	Recall $\left[\frac{N_{\text{correct detection}}}{N_{\text{tot injected}}} \right]$	-	0.967	1.000
TSNB detection	Precision $\frac{N_{\text{correct detection}}}{N_{\text{tot detection}}}$	-	1.000	1.000
	Recall $\frac{N_{\text{correct detection}}}{N_{\text{tot injected}}}$	-	1.000	1.000
Time	Processing time	-	75.1 s	71.4 s

Figure 4.6: Summary detection statistics comparison

The coverage already shows the qualitative difference predicted in Chapter 3: the coherence detector flags comprehensively the wide-band noise, even at higher SIR, while both perform equally good on the two narrow-band tones:

Detector mask coverage — mixed scenario

Mask	Coverage
Coherence: TSNB sub-mask	1.79%
Coherence: TSNB sub-mask	15.92%
Coherence: total mask (TSNB sub-mask OR TSNB sub-mask)	17.44%
Meyer HH	4.04%
Meyer VH	3.97%

Table 4.1: Fraction of the range-frequency/azimuth-time plane flagged by each detector on the mixed scenario.




Figure 4.7: SIR sweep detector comparison

4.4.2. Reconstruction and polarimetric performance




Figure 4.8: Visual comparison of the focused outputs for the mixed RFI injection.

Per-channel reconstruction vs. clean – pair HH/VH

Channel	Metric	Corrupted Input	Meyer+notch	Coh.+notch
HH	NMSE ($\rightarrow 0.196$)	0.381	0.329	0.215
	Δ SNR [dB] ($\uparrow$)	4.19	4.83	6.67
	Amplitude RMSE [dB] ($\downarrow$)	-7.57	-8.37	-10.12
	γ_{rec} ($\uparrow$)	0.828	0.817	0.884
	RFI-residual [%] ($\downarrow$)	100	86.2	56.4
VH	NMSE ($\rightarrow 0.196$)	1.162	0.367	0.306
	Δ SNR [dB] ($\uparrow$)	-0.65	4.35	5.15
	Amplitude RMSE [dB] ($\downarrow$)	-2.43	-8.05	-8.65
	γ_{rec} ($\uparrow$)	0.669	0.804	0.854
	RFI-residual [%] ($\downarrow$)	100	31.6	26.3

Table 4.2: Reconstruction of each channel against the clean reference. Arrows indicate the desired trend in the performance metric

Polarimetric coherence / covariance recovery – pair HH/VH

Metric	Clean	Corrupt. Input	Meyer+notch	Coh.+notch
Mean image coherence ($\rightarrow 0.091$)	0.091	0.285	0.105	0.097
Coherence-map RMSE vs. clean ($\downarrow$)	0	0.232	0.051	0.055
Unfoc. coherence at flagged cells ($\rightarrow 0.114$)	0.114	0.582	0.240	0.215
Cross-covariance rel. error ($\downarrow$)	0	5.565	0.693	0.712
Runtime [s]	—	—	168.2	152.3

Table 4.3: Recovery of the polarimetric pair statistics. Arrows indicate the desired trend in the performance metric.

As the presented results clearly show, the coherence method outperforms both qualitatively (looking at the coherence maps and amplitude comparisons) and quantitatively the reference method representing the current industry standard benchmark. Despite the results being promising, due to the limited amount of clean reference acquisitions available and the intrinsic limitations and uncertainties of simulations, the logical alternative is to proceed with real data testing in order to validate this method outside its tuning dataset.

4.4.3. Parameter sensitivity and tuning

The configuration frozen for the real-data application of Chapter 5 uses a frequency-averaging window $W = 10$ bins, the coherence detection threshold and line-length parameters of the detector as set in Chapter 3 (a coherence threshold of 0.62 and a minimum line requirement equal to twice the coherence computation window was found to be the best compromise between sensitivity and false alarm rejection). While the effect of the Threshold parameter and minimum line length on the algorithm sensitivity are clear from their definition, the tuning for the Window length parameter deserves a deeper discussion: while a shorter window for the coherence computation could theoretically increase the resolution capability and would allow the detector to track time-variable rfi structures more closely, as was already discussed in Chapter 3 a shorter window will bias the coherence estimate even when computed on a clean scene, increasing the baseline level from which the actual rfi spikes have to be isolated.

Following this reasoning, the theoretical false alarm rate previously computed can be confirmed through a threshold sweep during the synthetic noise injection simulation:




Figure 4.9: Empirical false alarm rate vs theoretical value; Empirical ROC curve and optimal Threshold

4.5. Discussion

The controlled experiments of this chapter establish the result that the rest of the thesis builds upon: while the coherence detector guarantees comparable removal under optimal conditions for the Meyer method [8], the proposed method outperforms the benchmark precisely in the regimes that matter for the Biomass mission, the cross-polarized channels characterized by a weaker signal and the heavily contaminated acquisitions, while requiring no background model and no prior knowledge of the interferer parameters. The simulation also exposes the intrinsic ceiling of notch-based removal: every correctly detected cell is a cell whose scene content is destroyed, so that beyond a certain contamination extent the reconstruction metrics of any detect-and-notch scheme are not able to recover the entire clean channel properties, no matter how accurate the mask. Overcoming this ceiling requires replacing the notch with a removal operator that subtracts the interference while preserving the underlying scene – the coherent cancellation approach developed and validated on real Biomass data in the next chapter.

5 | Application to Real Biomass Data

The simulation of Chapter 4 validated the detection stage against an exact ground truth and exposed the intrinsic limit of notch-based removal: each correctly flagged cell is a cell whose scene content is discarded along with the interference. This chapter moves to real Biomass commissioning-phase acquisitions, where no clean reference exists, and pursues two goals. The first is to replace the notch with a removal operator that subtracts the interference while preserving the underlying scene, exploiting the very cross-channel correlation that the detector uses to find the interference in the first place. The second is to assess the complete processing chain on real data through a four-method comparison that includes the official ESA Biomass algorithm as an external reference, on acquisitions independent from those used to tune the detector.

5.1. From Notching to Coherent Cancellation

5.1.1. The information loss of notching

Notching removes interference by zeroing the flagged cells (Equation (3.10)). Within a flagged cell the recorded value is the sum of scene, interference, and noise; setting it to zero removes the interference exactly but also discards the scene component $S(f, \tau)$ that shares the cell. For a sparse interferer this loss is negligible, but for a strong narrowband tone persisting over the whole azimuth extent, or a wideband burst spanning much of the band, the notch can excise a substantial portion of the range spectrum, leaving azimuth-stripe and ringing artefacts after refocusing. The simulation made this ceiling quantitative (Section 4.6), while on real data it would manifest as visible residual structure precisely where the contamination is densest.

5.1.2. The cancellation principle

The two channels of a receive pair offer a way out. As established in Section 3.1 and Equations (3.1)–(3.2), the cross-pol channel of a pair carries the same interference waveform as the co-pol channel (up to the time delay that can be neglected under our previous assumptions), while its scene component is statistically independent of the co-pol scene. The cross-pol channel is therefore an almost-ideal reference for a two-channel interference canceler (Section 1.8): it is strongly correlated with the interference to be removed since the scene power is lower in this channel, and uncorrelated with the signal to be preserved in the co-pol channel. For this reason, the main channel to be cleaned will be referred from now on as the Master channel (HH and VV in this case) and the channel used as reference for the RFI will be named Auxiliary (VH and HV respectively).

Applying the canceler to the H-receive pair, the cleaned co-pol channel at the flagged cells is

$$\hat{H}_{\text{HH}}(f, \tau) = H_{\text{HH}}(f, \tau) - W(f, \tau) H_{\text{VH}}(f, \tau), \quad (f, \tau) \in M, \quad (5.1)$$

where W is the cancellation weight and M is the detection mask of Chapter 3; cells outside M are left untouched. Under the assumptions that our signal satisfies, the weight that minimizes the residual coherence between the filtered signal and the auxiliary channel over a local window $\mathcal{W}$ is the local least-squares estimate over a window $\mathcal{W}$:

$$W(f, \tau) = \frac{\langle H_{\text{HH}} H_{\text{VH}}^* \rangle_{\mathcal{W}}}{\langle |H_{\text{VH}}|^2 \rangle_{\mathcal{W}}}, \quad (5.2)$$

This definition is also known as Wiener Filtering: where the interference dominates the reference, $W \cdot H_{\text{VH}}$ tends to r_H and the subtraction (5.1) removes r_H while leaving the independent scene term S_{HH} in place – exactly the cell content that the notch would have destroyed.

The cross-pol channel itself is not recoverable by the same method, because it provides its own reference and no third independent channel carrying the same interference is used for it. It is therefore notched, which is recovered as the degenerate case of (5.1) with unit weight applied to itself:

$$\hat{H}_{\text{VH}}(f, \tau) = H_{\text{VH}}(f, \tau) - 1 \cdot H_{\text{VH}}(f, \tau) = 0, \quad (f, \tau) \in M. \quad (5.3)$$

5.1.3. Local weight estimation

A unitary global weight would be optimal only if the two RFI realizations superposed to the received signals (sharing the same receive-polarity) were exactly identical. On the real case of the Biomass mission acquisitions this ideal hypothesis is not a valid assumption, as was proved by a series of experiments where the weight was set to $W = 1$ for the entire scene (presented in Chapter 5.3.2), so the weight has to be estimated locally, at the cost of an estimation variance that is related to the effective sample count of the chosen estimation window.

The implementation supports a spectrum of estimators of increasing sophistication – from a single global complex weight, through per-azimuth-column weights, to a local weight, windowed in both range and azimuth directions that adapts to non-stationary interference. After testing all these weight estimation options and observing the complex, highly non-stationary RFI environment observed in all the Biomass real data acquisitions, the best performing method that will be presented in the following section as representative of the coherent subtraction cancellation principle is the Wiener Filter with **Local-window Weight estimation**.

Two safeguards make the local estimate robust. First, the cancellation is gated by the binary detection mask: the full subtraction is applied only where the coherence detector has already declared a genuine RFI event, which prevents the canceler from acting on cells where the cross-pol channel is dominated by independent scene content and the weight estimate would be meaningless. Additionally, in order to reduce the discontinuity caused by a subtraction that is turned completely ON or OFF, a linearly scaled subtraction is applied on the pixels adjacent to the mask, allowing for a smooth transition between the filter application and the regions where the rfi power is deemed to low to apply a removal. This soft, locally adaptive form is the configuration carried into the benchmark of Section 5.2.

5.1.4. Strengths and weaknesses relative to notching

The cancellation method inherits the detection stage of the notch method unchanged – it operates on the same mask – so the two differ only in how flagged cells are treated, which makes their comparison a clean test of the removal step.

Strengths: Where the cancellation assumptions hold, it removes the interference while preserving the co-pol scene content of the flagged cells, recovering the information that the notch discards. The benefit grows with the contamination extent: the denser the

mask, the more scene cells the notch would have zeroed, and the larger the advantage of subtraction. Because it acts only inside the validated mask and is gated and soft-scaled, it does not degrade clean regions and does not introduce artifacts in regions where the estimation confidence is low.

Weaknesses: The method is limited by the cross-pol reference. Its effectiveness is bounded by how well the two channels' interference actually matches: any residual decorrelation between them – from the inter-channel acquisition delay of hypothesis H3 (Section 3.2), or from interferers whose envelope decorrelates faster than the pulse duration – leaves an unremovable residual and caps the achievable suppression depth. The local window for the weight estimation trades estimation variance against the ability to track $a(f, \tau)$: too small a window injects noise into the weight, too large a window would not track spatial or spectral variation. And the cross-pol channel, treated by (5.3), gains nothing over plain notching. The azimuth misalignment, in particular, is not a fixed property of the data but a correctable one, which directly motivates the improvement of Section 5.4. Thus, before continuing with this method development, it's necessary to verify that the operating hypothesis expressed in Chapter 3 are satisfied for the dataset we plan to process. The presence of near unity regions in the cross-channel coherence between two same-receive polarities, acquired over a natural scene, can confirm that at least some of the received interference is detectable and then removable by the proposed method. The following figure, representing the focused-domain coherence between the VV and HV polarity of one of the analyzed datasets confirms just that:




Figure 5.1: Cross-channel coherence (VV-HV) of the image processed in chapter 5.3, focused domain (t, τ) . The HH-VH image presents similar behaviour (See 5.6).

5.2. The Five-Method Comparison

5.2.1. Design of the experiment

The benchmark reported here was run on a Biomass commissioning-phase acquisition of a coastal region (independent from those used to set the detector parameters), chosen for its high range in amplitude that ensures a varied environment to test the method, with $N_r = 1528$ range samples and $N_a \approx 21000$ azimuth samples, spanning a surface of $\approx 4500 \text{ km}^2$, sampled at $f_s = 7.565 \text{ MHz}$ over a system bandwidth $B = 6 \text{ MHz}$. The scene is a strongly contaminated, mixed-class case: the coherence maps (Section 5.3) show a swept wide-band emitter (TVWB) concentrated in the early part of the aperture together with a population of persistent narrow-band tones and intermittent bursts (TSNB) spread across the rest of the swath, the combination flagging roughly a fourth of the range-frequency span. Both receive pairs were processed: the H-receive pair (HH cleaned through the VH reference) and the V-receive pair (VV cleaned through the HV reference) and the results are reported per receiving polarization (H and then V). One further scene was processed and the resulting comparison images are displayed in Appendix A.

Five methods were compared on the same tiles:

1. **Meyer detection + notching**: the single-channel reference detector (Section 2.5) followed by spectral notching, applied independently per channel;
2. **Coherence detection + notching**: the detector of Chapter 3 with the shared mask, followed by notching (Equation (3.10));
3. **Coherence detection + Simple coherent subtraction**: To isolate the contribution of the adaptive weight, a unit-weight variant of the cancellation is also reported (*Simple Subtraction*): the same coherence detector and shared mask as method 3 were used, but with the Wiener weight of Equation (5.2) replaced by $W \equiv 1$ in the flagged cells (the cross-pol channel is still notched). It subtracts the raw reference rather than its minimum-mean-square-error-scaled version, and shows the necessity for the introduction of the local weight estimation.
4. **Coherence detection + Wiener Weight Filtering**: the same detector, with the co-pol channel cleaned by the local Wiener cancellation of Equations (5.1)–(5.2) and the cross-pol channel notched (Equation (5.3));
5. **Official ESA Biomass filtering**: After the most effective of the four previous methods has been identified, an additional comparison has been carried against the algorithm currently deployed by the mission ground segment, used as an external

reference output.

In summary: methods 2, 3 and 4 share an identical detection stage and differ only in the removal step, isolating the contribution of cancellation; methods 1 and 2 share an identical removal step and differ only in detection, isolating the contribution of the coherence detector; method 5 gives an external reference for the method proposed in this work.

5.2.2. Evaluation without ground truth

On the real data three families of indirect, reference-free indicators were used in the final runs, all read comparatively across methods rather than as absolute figures of merit. The first is the *residual cross-channel coherence*, the cross-pol coherence magnitude γ measured *inside the originally flagged cells* (and, complementarily, over the whole focused image with a $W = 10$ window): a thorough removal returns the pair coherence from its interference-inflated value back toward the clean-scene baseline, so a *lower* post-correction γ indicates more complete suppression of the common interference. The second family characterizes the residual spectral structure left in the flagged range-frequency region, through three quantities: the *spectral-flatness ratio* (SFR), a dimensionless index of how much structured excess remains in the flagged band relative to the surrounding clean-band level, which should decrease toward its clean-baseline value when the interference excess is removed and grows when residual or newly injected structure is present; the *flagged-band dispersion* (σ_{Flag} , dB), the standard deviation of the range-power spectrum across the flagged bins, which falls as the flagged region is flattened toward the noise floor; and the *residual excess Power* (REP, dB), the excess power still standing in the flagged band above the interpolated clean-band floor, the most directly interpretable measure of leftover interference energy. The third family is qualitative: visual inspection of the focused amplitude images and of the range-frequency/azimuth time images after filtering for residual streaks, over-notching stripes and spectral holes. Because notching and cancellation act on the flagged cells in fundamentally different ways – the former zeros them, the latter refills them with estimates of the cleaned scene – the spectral-structure indices and the residual coherence must be read together: a method may drive the residual coherence down while leaving, or even worsening, the spectral structure, and only the combination distinguishes a genuine reconstruction from a destructive one.

5.2.3. Results




Figure 5.2: Results of the TSNB- and TVWB-windowed coherence computation, and their squared max value that better isolates the detected RFI




Figure 5.3: H-receive detection masks. Top row: the TSNB and TVWB coherence sub-masks; middle row: their union $m_{ref,H}$ (4.65% coverage), plus the Wiener mask carrying the detected coherence value. Bottom: the single-channel Meyer masks (3.72%/4.17%).

Detection mask coverage

Mask	Coverage
Coherence union $m_{\text{ref},H}$ (HH $\leftrightarrow$ VH)	4.6%
Coherence union $m_{\text{ref},V}$ (VV $\leftrightarrow$ HV)	4.8%
Meyer HH	3.72%
Meyer VH	4.17%
Meyer HV	3.68%
Meyer VV	3.27%

Table 5.1: Fraction of the range-frequency/azimuth-time plane flagged by each detector. The coherence detector produces a single, polarimetrically shared reference mask per receive pair; Meyer flags each channel independently.

Detection comparison Looking at the columns relative to the detection method proposed by Meyer et al. and the coherence-based detector paired with notching removal we can assess the detector performance alone: with the removal step held fixed at notching, replacing the Meyer detector by the coherence detector lowers the residual coherence at the flagged cells from 0.476 to 0.416 on HH–VH and from 0.478 to 0.429 on VV–HV, flagging the interference more completely and with a single shared mask for the pair rather than four individually computed per-channel masks (Table 5.1); as in simulation, the gain is clearest where the contamination is heaviest and the statistical nature of Meyer’s method operates in a suboptimal environment.

Reference-free spectral-structure metrics (lower is better)

Metric	Ch	Input	Meyer +notch	Coherence +notch	Simple subtraction ^a	Wiener
SFR	HH	0.654	0.449	0.449	0.988	0.511
σ_{Flag} [dB]	HH	3.64	2.44	2.66	4.60	2.81
REP [dB]	HH	13.06	5.39	5.95	15.49	7.20
SFR	VV	1.126	0.683	0.570	1.457	0.715
σ_{Flag} [dB]	VV	4.93	2.94	2.75	5.44	3.15
REP [dB]	VV	17.50	8.20	8.00	16.91	9.14

Table 5.2: Reference-free spectral-structure metrics on the real Data scene (lower is better). Only co-pol channels are shown. ^a unit-weight ($W \equiv 1$) subtraction.

Removal performances The comparison of the columns that share the same detection step (Coherence + Notch, Simple Subtraction, Wiener) allows to evaluate the removal effectiveness: the two subtraction-based removals reach the lowest residual coherence of all (Wiener 0.339 on HH–VH; Simple Subtraction 0.328 on VV–HV), well below the notch (0.416/0.429): actively cancelling the common reference decorrelates the pair more effectively than zeroing the flagged cells. Introducing in the comparison the other metrics allows to highlight one key difference: unit-weight subtraction reaches its low coherence distorting the image more than the initial RFI: its spectral-structure metrics are the worst in Table 5.2 and worse than the corrupted input itself (HH REP 15.49 dB vs. 13.06 dB at input; SFR 0.988 vs. 0.654), and the corresponding focused image (Figure 5.10, Simple Subtraction panel) shows the vertical interference columns *re-injected* more strongly than at input, because a fixed $W = 1$ over- or under-subtracts a reference that is neither perfectly scaled nor perfectly aligned to the interference. The local Wiener weight is then introduced to solve this issue: it manages a coherence reduction close to the notch level (0.339/0.336) while bringing the spectral-structure metrics back well below the input (HH REP 7.20 dB; SFR 0.511) and leaving the focused image free of both the original interference and any injected striping, at the cost of being the slowest method (576 s of execution time for $\approx 30 \cdot 10^6$ processed pixels).




Figure 5.4: Focused HH amplitude (dB): zoom on a clearly visible RFI artifact and the different method effects.




Figure 5.5: HH range-power spectra computed in the same azimuth range as 5.4: comparison between Input, Meyer method [8] and Wiener Weight filtering.

Scene Preservation The detector flags about a quarter of the range-frequency axis considering all flagged bins at some instant in azimuth (Table 5.1), so notching discards a corresponding fraction of the scene spectrum while the Wiener cancellation leaves in those cells the estimated cleaned scene. In the range-power spectra (Figure 5.12) the Wiener trace suppresses the interference humps while following the input spectral shape, and in the (f, τ) residual (Figure 5.18) it removes the tone streaks and the swept-band patch without leaving the zeroed cells of a notch: as expected, the advantage of coherent subtraction scales with the amount of mask coverage.

On the pure spectral-structure indices the notch is still the cleaner of the two valid removals (HH REP 5.4–5.9 dB for notching vs. 7.2 dB for Wiener), because zeroing a flagged cell removes all of its energy whereas a single block weight removes only the part of the reference that is correlated and correctly aligned with the interference. The residual that the Wiener weight leaves is the fraction of the interference that survives the azimuth misalignment of the receive pair (the one-pulse offset of hypothesis H3): the reference it subtracts is a time-shifted copy of the interference, and that shift can limit the achievable suppression depth.




Figure 5.6: Cross channel coherence (HH-VH) before and after every method's processing, focused domain (t, τ)




Figure 5.7: Cross channel coherence (VV-HV) before and after every method's processing, focused domain (t, τ)

Residual cross-channel coherence and runtime

	Input	Meyer +notch	Coherence +notch	Simple subtraction ^a	Wiener	ESA
<i>At flagged cells</i>						
<i>(γ in flagged f_r-τ region)</i>						
HH-VH	0.521	0.476	0.416	0.340	0.339	0.480
VV-HV	0.526	0.478	0.429	0.328	0.336	0.509
<i>Whole focused image</i>						
<i>(mean γ, $W = 10$)</i>						
HH-VH	0.291	0.206	0.169	0.165	0.155	0.273
VV-HV	0.298	0.188	0.171	0.144	0.160	0.277
Runtime [s]	—	537	511	482	576	n/a

Table 5.3: Residual cross-channel coherence (lower is better), measured both inside the originally flagged cells and over the whole focused image, and wall-clock runtime. ^a unit-weight subtraction.

Comparison with operational ESA filter On the focused amplitude (Figure 5.10) the ESA output is much smoother than any of the method that were compared, but the analysis of the focused time-domain and range-FFT amplitude after each method reveals that this method is preserving the least amount of useful information both in the spectral and in the time domain: the negative azimuth portion is left with broad zeroed regions in the spectrum and in the positive azimuth the time domain image is visibly artificially smoothed with a consequent loss of local information. Oppositely, the methods based on the coherence detector only act on the region that are influenced by the interferer, reducing the residual power in the spectral peaks while not affecting the region where the interference is negligible and spectral power is already near the clean base level. In particular the Wiener Weight removal method achieves similar visual performances in artifact removal while preserving the spectral characteristics even in the heavily RFI-corrupted frequency bins, and restoring the lowest level of cross-channel coherence after filtering. Inside the flagged cells (Table 5.3, Figure 5.8) the ESA output sits at an average coherence of 0.480 on HH-VH and 0.509 on VV-HV, only 8% and 3% below the corrupted input and above all the other compared methods, whereas the cross-channel Wiener Weight cancellation drives the same coherence to 0.339/0.336, a 35% reduction. Comparing the Wiener method and ESA algorithm using the spectrum (Figure 5.12) leads to similar results: ESA suppresses the interference humps but flattens

the range-power spectrum to a near-constant ~ 30 dB, reshaping the natural spectral roll-off, while the Wiener output follows the input shape and removes only the humps.




Figure 5.8: Residual cross-channel coherence γ_{TSNB} after cleaning (for channels HH–VH) in the unfocused domain; the mean over the flagged cells is printed in each panel.




Figure 5.9: Residual cross-channel coherence γ_{TSNB} after cleaning (for channels VV–HV) in the unfocused domain; the mean over the flagged cells is printed in each panel.

The following section is dedicated to the comparison figures between the proposed method and the current ESA algorithm for the co-pol channels, and for the detection effectiveness comparison between the coherence-based detection and the literature reference using the cross-pol channels (since for these the Wiener weight results are identical to the notch).

HH channel comparison figures:




Figure 5.10: Focused HH amplitude (dB) for the corrupted input and the five removals.




Figure 5.11: Range-power spectra of the HH channel: comparison between input, Wiener Weight method and Current ESA Filter




Figure 5.12: HH channel comparison: Focused domain range-FFT, showing residual RFI power and eventual notching traces.

VH Channel comparison figures:




Figure 5.13: Focused VH amplitude (dB) for the corrupted input and the five removals.




Figure 5.14: Range-power spectra of the VH channel: comparison between input, Meyer detection, Coherence-based detection and Current ESA Filter

VV channel comparison figures:




Figure 5.15: Focused VV amplitude (dB) for the corrupted input and the five removals.




Figure 5.16: Range-power spectra of the VV channel: comparison between input, Wiener Weight method and Current ESA Filter

HV channel comparison figures:




Figure 5.17: Focused HV amplitude (dB) for the corrupted input and the five removals.




Figure 5.18: Range-power spectra of the HV channel: comparison between input, Meyer detection, Coherence-based detection and Current ESA Filter

5.3. Interference Environment of the Analyzed Scene

On this specific acquisition the detector output reads as a direct map of the local P-band environment as well (Figure 5.19), allowing us to conduct a preliminary Interference Environment Analysis using this data. In the H-receive channels, a swept wide-band emitter (TVWB) appears as a coherent diagonal band near -2 to -3 MHz of range frequency, confined to the first part of the aperture (azimuth $-8 \cdot 10^4$ to $-3 \cdot 10^4$), its slope marking the frequency-time sweep of the source. The rest of the plane is populated by narrow-band contributions (TSNB): a set of persistent tones that draw nearly horizontal high-coherence streaks across azimuth, and a denser cluster of compact, azimuth-localized bursts concentrated in the second half of the swath (azimuth $+2 \cdot 10^4$ to $+8 \cdot 10^4$). The contamination spans essentially the whole ± 3 MHz band, and the whole azimuth extension is affected, with two particularly dense regions just described; The V-receive channels do not show any substantial wide-band detection, but highlight two different interferers categories as well: most of the persistent detections are distinctively stationary in frequency and narrow band, but the narrow-band bursts seen in the previous mask affect this polarization as well. Together the two reference masks flag 4.6% of the plane on the H-receive pair and 4.8% on the V-receive pair (Figure 5.3, Table 5.1), against only $\sim 4\%$ flagged per channel by the single-channel Meyer detector, which captures the strong discrete tones but misses some of the diffuse swept and burst energy.

This picture is consistent with the geographic expectation of Section 2.1 for a densely populated site: a persistent narrow-band population of the kind produced by terrestrial fixed-service and TETRA-type emitters, overlaid with intermittent wide-band bursts of radar-type origin. It also illustrates the motivation for a data-driven detector: the burst clusters are diffuse, time-varying features that may not be detectable during pre-launch surveys, yet they are flagged here directly from the cross-channel coherence without any prior knowledge on the interferer parameters.




Figure 5.19: H-receive (HH/VH) detection map for the coherence detector in the range-frequency / azimuth-time plane.




Figure 5.20: V-receive (VV/HV) detection map for the coherence detector in the range-frequency / azimuth-time plane.

5.4. Possible Improvements

5.4.1. Azimuth alignment of the receive pair

The single most direct improvement follows from hypothesis H3 of Section 3.2. The two channels of a receive pair are acquired one pulse apart (Equation (3.3)), so corresponding azimuth samples of H_{HH} and H_{VH} are offset by $\Delta\tau = 1/\text{PRF}$. The detector tolerates this offset because it only needs the coherence magnitude to rise, but the cancellation of Equation (5.1) is degraded by it: the reference it subtracts is a time-shifted copy of the interference, and the residual misalignment limits the suppression depth.

If the system timing is known accurately, the offset can be compensated before the cancellation by shifting the reference channel in azimuth by the corresponding integer number of columns,

$$\hat{H}_{HH}(f, \tau) = H_{HH}(f, \tau) - W H_{VH}(f, \tau + \Delta\tau), \quad (f, \tau) \in M, \quad (5.4)$$

re-aligning the two channels to the same effective acquisition time before the subtraction. For a standard interleaved quad-pol product, one azimuth column of each channel corresponds to $1/(\text{PRF}/2)$ seconds, and the shift required is the integer number of columns nearest to $\Delta\tau$. The improvement is expected to be largest for the interferers whose envelope decorrelates near the pulse scale – exactly those for which the un-shifted cancellation leaves the largest residual – and negligible for those interference sources that present coherence close to unity even after the reception delay.

On the dataset results presented in Section 5.2 this misalignment is the leading explanation for the gap between the two valid removals in Table 5.2: the Wiener cancellation matches notching on residual coherence but leaves larger residual peaks in the spectrum plot

(~ 10 dB vs. ~ 5 – 6 dB on HH), exactly the signature of a reference subtracted with the wrong effective acquisition time. The integer azimuth shift required to re-align the pair follows from the product timing as the number of columns nearest to $\Delta\tau = 1/\text{PRF}$; with the azimuth column spacing $1/(\text{PRF}/2)$ of the interleaved quad-pol product, this is exactly 1 azimuth column if the resulting image is sampled exactly at this rate; If during processing the azimuth spacing is modified from its original sampling rate this relation is not valid anymore, and the necessary shift would become a sub-unity value, necessitating more complex techniques to recover the higher interferer coherence that would improve noticeably both the detector and coherent subtraction removal method. The shifted-cancellation variant was not run in the present comparison; the expected effect is a reduction of the residual hump toward the notch level for the interferers which decorrelate below the detection threshold faster than $\Delta\tau = 1/\text{PRF}$, with negligible change for those already coherent near unity after the reception delay.

5.5. Discussion on Real Data Results

The real Biomass data application allows us to evaluate the coherence detector generalization capability to acquisitions it was not tuned on, and quantify the recovered scene content by the subtraction method that notching discards, with the largest gains where contamination is densest – the regime in which notching is weakest. Additionally, compared against the operational ESA algorithm, the proposed method implementing Wiener weight filtering matches it in visual image quality while roughly halving the residual cross-channel coherence at the flagged cells (a 35% reduction against ESA’s $\sim 5\%$) and preserving the scene spectral signature that the ESA algorithm reshapes. The principal residual limitation, the inter-channel azimuth offset, is not intrinsic but correctable given accurate system timing, pointing to the alignment refinement of Section 5.4.1 as the natural next step. Beyond the scene analyzed in detail here, one more dataset was processed with the same coherent-subtraction method and compared in the same way; the results indicate that the method’s working hypotheses are satisfied to a lesser degree, with the detector flagging only a subset of the interferers present. The full comparison figures for these acquisitions are collected in Appendix A.

5.5.1. Comparison with the Literature

It is useful to compare this work’s proposed method against the methods reviewed in Section 2.5, restricting the comparison to the common performance metrics – detection probability and false-alarm rate against ground truth, reconstruction coherence and SNR

improvement against the clean reference, and residual interference power.

Versus single-channel spectral detection (Meyer et al., 2013 [8]). Measured with the reconstruction metrics of Chapter 4, the coherence detector reconstructs each channel better than the single-channel reference on both polarizations: on the co-pol channel it reaches $\text{NMSE} = 0.22$ against Meyer’s 0.33, and on the cross-pol channel 0.31 against 0.37, with the additional advantages of needing no background model, having cross-channel support for the weak channel and polarimetric mask consistency, while on the joint polarimetric statistics the two are comparable for narrow-band tones. The surviving interference is likewise cut to 56% and 26% of the input against the single-channel 86% and 32%. The same dominance appears in detection: across the SIR sweep its P_d exceeds Meyer’s at every severity on both channels, the two curves separating most on the cross-pol channel – near 0 dB SIR – where the single-channel test saturates close to 0.5, while the coherence detector’s false-alarm rate falls with increasing interference rather than rising. This advantage becomes decisive in the cross-pol low-SNR regime, where the single-channel background estimate degrades.

Versus cross-polarization coherence detection (Uozumi et al., 2023) [14]. The present detector follows the same physical intuition but extends the original formulation from a coherence indicator to a complete, validated mask-construction pipeline: oriented estimators for the two RFI classes, a threshold tied to the estimator’s null statistics, morphological line detection and a false-alarm rejection stage, while discarding the focused-domain coherence detection that was found to degrade performance when implemented.

Versus sidelobe-cancellation suppression (Wang et al.) [15]. Coherent cancellation using one polarimetric channel as the reference for another is known to outperform notching when the inter-channel interference is coherent, but its performance is limited by the use of a single global weight, whereas the optimal weight varies with frequency and azimuth. The removal developed in Section 5.1 addresses this with a locally estimated weight, applied only on the validated detection mask, so that cancellation acts only where a genuine interference event has been confirmed and tracks the spatially varying coupling.

Versus the operational ESA algorithm. As an external, end-to-end reference, the deployed ESA filter produces clean, smooth imagery on a par with the Wiener output and free of the input’s interference columns; on visual quality the two are comparable. The proposed method exceeds it, however, on the residual cross-channel coherence at the flagged cells (0.34/0.33 for the Wiener cancellation against 0.48/0.51 for ESA on the two receive pairs, Figure 5.8, Figure 5.9). The two also differ spectrally: ESA whitens the range-power spectrum to a near-constant level, removing the interference humps but

distorting the scene’s natural spectral shape, while all the coherence-based cancellation methods suppress the spectral peaks yet follow the input shape. The chain thus matches the operational baseline where it is strongest (image cleanliness) and exceeds it where the cross-channel physics gives it leverage (residual coherence and spectral-shape preservation), at the cost of a higher per-scene runtime.

5.5.2. Categorization of the Detected Interferers

Analysis of the coherence detection maps across all the available acquisitions reveals a taxonomy of RFI morphologies that can be organised into the following categories, based on their signatures in the range-frequency / azimuth-time domain.

Persistent Narrow-band Tones (TSNB, Category I)

The most frequently observed RFI type consists of spectrally stable, azimuth-persistent narrow-band tones. These appear as horizontal stripes of high coherence at fixed frequency positions in the (f, τ) domain, extending across the full azimuth extent of the processed tile. The spectral width of each stripe is consistent with the bandwidth of a narrow-band carrier (a few to tens of kHz – sub-bin at the Biomass range-frequency resolution). The dominant sources of this category are believed to be terrestrial professional mobile radio networks (TETRA) and similar fixed-frequency communication systems operating in or near the 435 MHz allocation. This category is reliably detected by the TSNB coherence estimator and completely suppressed by both the notch and the Wiener-validated filter.

Wide-band Burst Interference (TVWB, Category II)

Wide-band burst interference occupies an extended frequency range (typically tens of percent of the 6 MHz system bandwidth) for a limited azimuth duration, appearing as vertical or near-vertical elevated-coherence features in the (f, τ) domain. The azimuth extent of detected bursts ranges from a few pulses to several hundred azimuth lines, consistent with pulse trains from ATC or maritime surveillance radars that illuminate the Biomass footprint for a limited orbital arc. This category is detected by the TVWB estimator and is more challenging for spectral notching: because the contamination spans many frequency bins, the bandwidth loss from notching is more significant, and the Wiener-validated method offers a more visible advantage in preserving range resolution.

Frequency-Drifting Wide- and Narrow-band Tones (Category III)

A subset of detections shows slow frequency drift along the azimuth axis: the spectral position of the coherence stripe shifts gradually across longer time spans, and can involve either isolated frequencies or wider bands. This behaviour is consistent with Doppler-shifted carrier transmissions from moving platforms (vehicles, airborne systems) or with thermally unstable oscillators in ground-based emitters. The TSNB estimator detects this category very reliably for drift rates slower than approximately one frequency bin per W_F azimuth lines; faster drift rates are partially captured by the TVWB estimator. Detection is generally complete, but the resulting mask has a slightly irregular horizontal extent – the line-length filter must be tuned accordingly to accommodate the short segments of high coherence that correspond to the fast time-frequency transitions between drift segments.

Burst Narrow-band RFI (Category IV)

In some acquisitions, narrow-band tones are present only for short azimuth extents, producing a characteristic localized high-power spot. These signatures are detected by both detector orientations, and the OR-combination of the two masks handles the category correctly; the residual power after removal can still be visible in the post-filtering coherence maps, even when the corresponding artifacts in the focused time domain are completely suppressed.

On the acquisition analyzed in Section 5.2 the masks resolve two coexisting classes: a swept wide-band emitter (Cat. II–III), appearing as a coherent diagonal band near -2 to -3 MHz of the H-polarity receive mask and confined to the early part of the aperture, and a narrow-band population (Cat. I) of persistent tones – nearly horizontal high-coherence streaks – which dominate the V-polarity receive mask, together with sparse intermittent bursts (Cat. IV) concentrated in the second half of the swath for both (Figure 5.19 and 5.20).

Scene	Dominant class	Secondary	Coverage	
			H-pair	V-pair
Section 5.2	TSNB tones + bursts	TVWB swept band	4.6%	4.8%
Appendix A scene	TNSB tones	n/a	1.1 %	0.47 %

Table 5.4: Categorisation of the coherent interferers detected across the analysed Biomass acquisitions. Coverage is the fraction of the range-frequency/azimuth-time plane flagged by the detection mask of each receive pair.

6 | Conclusions and Future Developments

The Biomass mission was conceived to answer a precise scientific need: above-ground biomass is one of the largest estimated terms (and the single one associated with the highest uncertainty) in the global carbon cycle, and mapping it globally means reaching the woody structures that store most of it – something only the long P-Band wavelength can achieve from space. That same choice carries an intrinsic trade-off: the P-Band radar operates in a narrow allocation shared with a dense population of terrestrial transmitters, so radio frequency interference is a frequent feature of its data: left uncorrected it raises the noise floor, biases the calibrated backscatter, and erodes the coherence on which the mission’s interferometric and tomographic products depend, with the weak cross-polarized channels suffering most. Because the geographical distribution of P-Band interferers is largely undocumented, any mitigation tied to a pre-compiled catalogue is unviable, while the single-channel, background-based methods that have been explored for earlier missions lose accuracy precisely in the heavily contaminated conditions Biomass must routinely handle. A mitigation built around the structure of the data itself was therefore needed, and providing one is what this thesis set out to do.

The problem was approached in two stages, detection and suppression, both addressed through the same physical property: the coupling that the shared receive port imposes between the channels of a polarimetric pair. The two methods that resulted – a catalogue-free coherence detector and a scene-preserving coherent cancellation – were developed to operate as one chain, with the mask produced by the first becoming the support on which the second acts.

Taken together, the controlled-simulation and real-data assessments of the preceding chapters support a consistent reading of what that chain achieves. Its effectiveness is concentrated exactly where amplitude- and power-based methods are weakest: coherence-based detection is highly effective on the low-backscatter cross-polarized channels, and the coherent cancellation method operates best in the densely contaminated acquisitions (where

the target signal to be removed is dominating the scene). Additionally, because the proposed suppression stage cancels the interference rather than zeroing the affected frequency bins, it returns a cleaned product that retains the scene content a spectral notch would have discarded – the property that matters most for the interferometric and tomographic uses for which Biomass exists. The method is in this sense not merely an alternative to the established baseline but a complement to it, strongest in the regimes the baseline handles least well.

These same results also mark the boundaries of the present implementation, and it is from those boundaries that the most useful next steps are drawn. Some follow from adopted hypotheses that can now be relaxed, while others from design choices kept deliberately narrow in order to validate the core idea before generalizing it: the following section collects the most promising of these directions.

6.1. Future Developments

Azimuth alignment of the receive pair. The inter-channel acquisition delay (hypothesis H3, Section 3.2) caps both the observable coherence and the cancellation depth. The Wiener output residual coherence is always comparable to or better than the notch’s, yet it leaves a larger residual peak in most cases (~ 10 dB against the notch’s ~ 5 – 6 dB on HH) – the signature of a correctly detected interference subtracted at the wrong effective acquisition time. Compensating the one-pulse offset before cancellation is expected to close most of that gap for the slowly decorrelating interferers, bringing the residual toward the notch level while preserving the scene-preservation advantage. The variant was not run in the present comparison, so this remains a theoretically-backed expectation.

Full exploitation of both receive pairs in removal. The cancellation currently uses a single auxiliary channel for the Wiener estimation, yet in the same dataset a similar coherence coupling was detected in the opposite-receive co-pol/cross-pol pairs (the HH/HV high-coherence regions partially overlap those found with the HH/VH pair). A joint formulation that uses multiple auxiliary channels as corrupted references could improve both detection and clean-signal recovery.

Operational considerations. The pipeline reduces to separable convolutions and connected-component analysis, scaling linearly with the number of pixels, which makes it attractive for large-volume ground-segment processing. Across the processed datasets the proposed solution is slower than the single-channel method on every individual scene, but with a clear trend toward higher efficiency at larger pixel counts: the measured runtimes (Table 6.1) fit straight lines whose Wiener slope, $14.4 \mu\text{s}/\text{pixel}$, is well below the single-channel

one, $20.7 \mu\text{s}/\text{pixel}$, so the two fitted lines cross at $N^* \approx 39 \text{ Mpixel}$ (Figure 6.1). On the smaller products the cancellation is the slower of the two, but beyond N^* it becomes the more efficient pipeline. The limited number of samples makes the crossover point itself highly uncertain, but the difference in per-pixel cost is large enough to make the favourable large-scale trend a reasonable expectation. Additionally, the negative intercept resulting from the Meyer method runtimes interpolation is unrealistic, but can be explained by the low number of samples, together with the lower linear fitting of the computation time for this method, which is visible as non linear even within Figure 6.1.

Runtime scaling of the two detection and removal pipelines

	N [Mpix]	Meyer+notch	Wiener weight filtering
<i>Measured runtime [s]</i>			
Appendix B	17	271	388
Appendix A	23	344	485
Chapter 5.2	30	537	576
<i>Least-squares fit $t = a + bN$</i>			
slope b [$\mu\text{s}/\text{pixel}$]	—	20.7	14.4
intercept a [s]	—	−98	147
R^2	—	0.92	0.996

Table 6.1: Measured per-scene runtimes of the single-channel (Meyer+notch) and the local Wiener-cancellation pipelines, and their linear fits versus total image size N .




Figure 6.1: Runtime interpolation and prediction.

6.2. Closing Remark

This work shows that the polarimetric structure of the Biomass data is not merely an end product to be protected from interference but also the means to detect and suppress it: the same shared-receive coupling that makes the cross-pol channels especially vulnerable to RFI is what makes one channel an effective reference for cleaning another.


One additional consideration on the content of this thesis is the timely arrival of its contribution: BIOMASS is the first synthetic aperture radar to operate at P-band from space, and until its 2025 in-orbit commissioning no spaceborne P-band data existed: the established RFI detection and suppression literature was developed either on airborne P-band campaigns or on space-based acquisitions at other frequencies, referring to previous missions. The algorithm presented here is therefore among the first methods reported in the open literature to detect and suppress radio frequency interference directly on real spaceborne P-band data, alongside the mission's own operational processing baseline: arriving this early in the mission's life, the present study is offered in the hope that it may play some part in strengthening BIOMASS's resilience to interference, and consequently improving the reliability of the forest and carbon measurements the mission exists to deliver.

Bibliography


- [1] R. Bamler and P. Hartl. Synthetic aperture radar interferometry. *Inverse Problems*, 14(4):R1–R54, 1998. doi: 10.1088/0266-5611/14/4/001.
- [2] I. G. Cumming and F. H. Wong. *Digital Processing of Synthetic Aperture Radar Data: Algorithms and Implementation*. Artech House, Norwood, MA, USA, 2005.
- [3] S. Haykin. *Adaptive Filter Theory*. Prentice Hall, Upper Saddle River, NJ, USA, 4th edition, 2002.
- [4] International Telecommunication Union. Feasibility of sharing between active spaceborne sensors and other services in the range 420–470 MHz. Technical Report Recommendation ITU-R RS.1260-2, ITU Radiocommunication Sector (ITU-R), Geneva, Switzerland, 2017.
- [5] J.-S. Lee and E. Pottier. *Polarimetric Radar Imaging: From Basics to Applications*. CRC Press, Boca Raton, FL, USA, 2009.
- [6] R. T. Lord and M. R. Inggs. Efficient RFI suppression in SAR using a LMS adaptive filter integrated with the range/Doppler algorithm. In *Electronics Letters*, volume 35, pages 629–630, 1999. doi: 10.1049/el:19990437.
- [7] Y. Mao, Y. Huang, X. Yu, Y. Wang, M. Tao, Z. Zhang, Y. Yang, and W. Hong. Radio frequency interference mitigation in SAR systems via multi-polarization framework. *IEEE Transactions on Geoscience and Remote Sensing*, 62:5210216, 2024. doi: 10.1109/TGRS.2024.3384953.
- [8] F. J. Meyer, J. B. Nicoll, and A. P. Doulgeris. Correction and characterization of radio frequency interference signatures in L-band synthetic aperture radar data. *IEEE Transactions on Geoscience and Remote Sensing*, 51(10):4961–4972, 2013. doi: 10.1109/TGRS.2013.2252469.
- [9] S. Quegan, T. Le Toan, J. Chave, J. Dall, J.-F. Exbrayat, D. Ho Tong Minh, M. Lomas, M. Mariotti D’Alessandro, P. Paillou, K. Papathanassiou, F. Rocca, S. Saatchi, K. Scipal, H. Shugart, T. L. Smallman, M. J. Soja, S. Tebaldini, L. Ulander, L. Vil-

- lard, and M. Williams. The European Space Agency BIOMASS mission: Measuring forest above-ground biomass from space. *Remote Sensing of Environment*, 227:44–60, 2019. doi: 10.1016/j.rse.2019.03.032.
- [10] M. Tao, J. Su, Y. Huang, and L. Wang. Mitigation of radio frequency interference in synthetic aperture radar data: Current status and future trends. *Remote Sensing*, 11(20):2438, 2019. doi: 10.3390/rs11202438.
- [11] S. Tebaldini. Single and multipolarimetric SAR tomography of forested areas: A parametric approach. *IEEE Transactions on Geoscience and Remote Sensing*, 48(5): 2375–2387, 2010. doi: 10.1109/TGRS.2009.2037748.
- [12] S. Tebaldini. Remote sensing for earth observation and surveillance – lecture notes. Course material, Politecnico di Milano, 2025. M.Sc. in Space Engineering, Politecnico di Milano.
- [13] R. Touzi, A. Lopès, J. Bruniquel, and P. W. Vachon. Coherence estimation for SAR imagery. *IEEE Transactions on Geoscience and Remote Sensing*, 37(1):135–149, Jan. 1999. doi: 10.1109/36.739146.
- [14] A. Uozumi, A. Hirose, and R. Natsuaki. RFI detection using cross-polarization coherence for dual- and full-polarimetric synthetic aperture radar. In *Proceedings of the URSI General Assembly and Scientific Symposium (GASS)*, Sapporo, Japan, 2023.
- [15] J. Wang, L. Zhang, Z. Meng, S. Wei, and L.-L. Chen. General RFI suppression with sidelobe cancellation filtering for dual-polarization SAR images. *IEEE Geoscience and Remote Sensing Letters*, 20:4012005, 2023. doi: 10.1109/LGRS.2023.3321321.
- [16] B. Widrow, J. R. Glover, J. M. McCool, J. Kaunitz, C. S. Williams, R. H. Hearn, J. R. Zeidler, E. Dong, and R. C. Goodlin. Adaptive noise cancelling: Principles and applications. *Proceedings of the IEEE*, 63(12):1692–1716, 1975. doi: 10.1109/PROC.1975.10036.


A | Appendix A




(a) HH-VH detection mask




(b) VV-HV detection mask




(a) HH-VH residual coherence (focused domain)




(b) HH amplitude comparison




(c) VH amplitude comparison




(a) HH-VH residual TSNB coherence (defocused domain)




(b) HH average spectrum comparison




(c) VH average spectrum comparison






(a) VV-HV residual coherence (focused domain)



(b) VV amplitude comparison



(c) HV amplitude comparison



(c) HV average spectrum

Figure A.6: Comparison figure list of a desert scene (exact location redacted)

List of Figures

3.1	Double estimation window working principle	32
4.1	Architecture of the controlled validation environment. The injection takes place in the raw-image domain, where physical interference actually enters the receiver	40
4.2	repeated defocus-refocus effect and introduced distortions	41
4.3	example of the expected behaviour of the artificial RFI injection	42
4.4	Simulated RFI for the mixed case injected over a blank image background	43
4.5	RFI Injection and detection comparison: coherence elevates toward unity exactly over the injected structures, as predicted by Equation (3.5).	44
4.6	Summary detection statistics comparison	48
4.7	SIR sweep detector comparison	49
4.8	Visual comparison of the focused outputs for the mixed RFI injection.	49
4.9	Empirical false alarm rate vs theoretical value; Empirical ROC curve and optimal Threshold	51
5.1	Cross-channel coherence (VV-HV) of the image processed in chapter 5.3, focused domain (t, τ) . The HH-VH image presents similar behaviour (See 5.6).	56
5.2	Results of the TSNB- and TVWB-windowed coherence computation, and their squared max value that better isolates the detected RFI	59
5.3	H-receive detection masks. Top row: the TSNB and TVWB coherence sub-masks; middle row: their union $m_{\text{ref},H}$ (4.65% coverage), plus the Wiener mask carrying the detected coherence value. Bottom: the single-channel Meyer masks (3.72%/4.17%).	59
5.4	Focused HH amplitude (dB): zoom on a clearly visible RFI artifact and the different method effects.	61
5.5	HH range-power spectra computed in the same azimuth range as 5.4: comparison between Input, Meyer method [8] and Wiener Weight filtering.	62

5.6	Cross channel coherence (HH-VH) before and after every method's processing, focused domain (t, τ)	63
5.7	Cross channel coherence (VV-HV) before and after every method's processing, focused domain (t, τ)	63
5.8	Residual cross-channel coherence γ_{TSNB} after cleaning (for channels HH-VH) in the unfocused domain; the mean over the flagged cells is printed in each panel.	65
5.9	Residual cross-channel coherence γ_{TSNB} after cleaning (for channels VV-HV) in the unfocused domain; the mean over the flagged cells is printed in each panel.	65
5.10	Focused HH amplitude (dB) for the corrupted input and the five removals.	66
5.11	Range-power spectra of the HH channel: comparison between input, Wiener Weight method and Current ESA Filter	66
5.12	HH channel comparison: Focused domain range-FFT, showing residual RFI power and eventual notching traces.	67
5.13	Focused VH amplitude (dB) for the corrupted input and the five removals.	67
5.14	Range-power spectra of the VH channel: comparison between input, Meyer detection, Coherence-based detection and Current ESA Filter	68
5.15	Focused VV amplitude (dB) for the corrupted input and the five removals.	68
5.16	Range-power spectra of the VV channel: comparison between input, Wiener Weight method and Current ESA Filter	69
5.17	Focused HV amplitude (dB) for the corrupted input and the five removals.	69
5.18	Range-power spectra of the HV channel: comparison between input, Meyer detection, Coherence-based detection and Current ESA Filter	69
5.19	H-receive (HH/VH) detection map for the coherence detector in the range-frequency / azimuth-time plane.	70
5.20	V-receive (VV/HV) detection map for the coherence detector in the range-frequency / azimuth-time plane.	71
6.1	Runtime interpolation and prediction.	79
A.6	Comparison figure list of a desert scene (exact location redacted)	88

List of Tables

1	Radar frequency bands for Earth observation.	2
3.1	Structural comparison between the reference single-channel detector and the proposed coherence-based detector.	37
4.1	Fraction of the range-frequency/azimuth-time plane flagged by each detector on the mixed scenario.	48
4.2	Reconstruction of each channel against the clean reference. Arrows indicate the desired trend in the performance metric	50
4.3	Recovery of the polarimetric pair statistics. Arrows indicate the desired trend in the performance metric.	50
5.1	Fraction of the range-frequency/azimuth-time plane flagged by each detector. The coherence detector produces a single, polarimetrically shared reference mask per receive pair; Meyer flags each channel independently. . .	60
5.2	Reference-free spectral-structure metrics on the real Data scene (lower is better). Only co-pol channels are shown. ^a unit-weight ($W \equiv 1$) subtraction.	60
5.3	Residual cross-channel coherence (lower is better), measured both inside the originally flagged cells and over the whole focused image, and wall-clock runtime. ^a unit-weight subtraction.	64
5.4	Categorisation of the coherent interferers detected across the analysed Biomass acquisitions. Coverage is the fraction of the range-frequency/azimuth-time plane flagged by the detection mask of each receive pair.	75
6.1	Measured per-scene runtimes of the single-channel (Meyer+notch) and the local Wiener-cancellation pipelines, and their linear fits versus total image size N	79

List of Symbols

Symbol	Description	Unit
<i>SAR geometry and signal model</i>		
t	Fast time (within a pulse)	s
τ	Slow time / azimuth time	s
f_r	Range frequency	Hz
c	Speed of light	m/s
λ	Radar wavelength	m
f_0	Carrier frequency	Hz
f_s	Range sampling rate	Hz
θ	Incidence angle	rad
R, R_0	Slant range / broadside slant range	m
r	Slant-range coordinate	m
x, y	Azimuth / ground-range coordinate	m
v	Platform velocity	m/s
L_x	Physical antenna length (azimuth)	m
T_p	Transmitted pulse duration	s
B	Transmitted (system) bandwidth	Hz
K_r	Chirp rate	Hz/s
PRI	Pulse repetition interval	s
PRF	Pulse repetition frequency	Hz
δ_r	Slant-range resolution	m
δ_y	Ground-range resolution	m
δ_x	Azimuth resolution	m
$g(t)$	Transmitted chirp (LFM) envelope	—
$h(t, \tau)$	Baseband received data	—
$d_{rc}(t, \tau)$	Range-compressed pulse	—
$I(x, r)$	Focused complex image	—

continued on next page

Symbol	Description	Unit
N_r, N_a	Number of range / azimuth samples	–
P_{tx}, P_{rx}	Transmitted / received power	W
G	Antenna gain	–
σ^0	Backscatter coefficient (normalised RCS)	–
σ_{NESZ}^0	Noise-equivalent sigma zero (NESZ)	dB
$n(t, \tau)$	Thermal noise	–
K_B	Boltzmann constant	J/K
T_{sys}	System noise temperature	K
L_{eq}	Equivalent number of looks	–
<i>Polarimetry, coherence and tomography</i>		
S	2×2 complex scattering matrix	–
$S_{HH}, S_{HV}, S_{VH}, S_{VV}$	Scattering-matrix / channel components	–
C	Polarimetric covariance matrix	–
I_{pq}	Focused polarimetric channel, $pq \in \{HH, HV, VH, VV\}$	–
$\gamma, \gamma $	Coherence (complex) and its magnitude	–
$\hat{\gamma}$	Sample (estimated) coherence	–
ℓ_{coh}	Number of looks in a coherence estimate	–
ϕ	Interferometric phase	rad
σ_ϕ^2	Interferometric phase-error variance	rad ²
k_z	Height-to-phase (vertical wavenumber) factor	rad/m
$b_\perp, b_n$	Perpendicular / n -th orbital baseline	m
Δb	Total baseline span	m
δ_v	Elevation (tomographic) resolution	m
$W_{\text{opt}}(f)$	Optimal Wiener (MMSE) filter response	–
S_{ys}, S_{yy}	Cross- / auto-power spectra	–
<i>Radio frequency interference</i>		
(f, τ)	Range-frequency / azimuth-time domain	(Hz, s)
$r(t, \tau)$	Interference term	–
r_{NB}, r_{WB}	Narrowband / wideband interference	–
ν_0	Narrowband interferer baseband frequency	Hz
A_0	Interference amplitude	–
$L_{\text{az}}(\tau)$	Azimuth activity pattern, $\in \{0, 1\}$	–

continued on next page

Symbol	Description	Unit
T_d	Wideband burst duration	s
B_{wb}	Wideband interference bandwidth	Hz
f_c	Burst centre frequency	Hz
SIR	Signal-to-interference ratio	dB
P_s	Mean scene power	W
P_r	Mean interference power	W
<i>Coherence-based detector</i>		
$H_{pq}(f, \tau)$	Channel data in the (f, τ) domain	—
H_{co}, H_x	Co- / cross-polarised channel of a receive pair	—
$\hat{\gamma}_{\text{TSNB}}$	Azimuth-oriented coherence estimator	—
$\hat{\gamma}_{\text{TVWB}}$	Frequency-oriented coherence estimator	—
W	Coherence estimation window length	samples
W_F	TSNB window length (azimuth lines)	—
W_T	TVWB window width (frequency bins)	—
γ_{th}	Coherence detection threshold	—
$\ell_{\min}$	Minimum line length	samples
$M(f, \tau)$	Binary detection mask	—
$\Delta\tau$	Inter-channel acquisition delay	s
τ_c^{RFI}	Interference coherence time	s
<i>Coherent cancellation (removal)</i>		
$\hat{H}_{HH}$	Cleaned co-pol (master) channel	—
$W(f, \tau)$	Local Wiener cancellation weight	—
H^{notch}	Notched data	—
N	Total image size	pixel
<i>Performance metrics</i>		
M_{gt}	Ground-truth detection mask	—
P_d	Detection probability	—
P_{fa}	False-alarm rate	—
NMSE	Normalised mean-square error	—
ΔSNR	SNR improvement	dB
γ_{rec}	Reconstruction coherence	—
C_{ab}	Off-diagonal polarimetric covariance element	—
ε_C	Cross-covariance relative error	—

continued on next page

Symbol	Description	Unit
SFR	Spectral-flatness ratio	—
σ_{Flag}	Flagged-band spectral dispersion	dB
REP	Residual excess power	dB

List of Acronyms

AGB	Above-Ground Biomass
ATC	Air-Traffic Control
CCG	Complex Circular Gaussian
EES	Earth Exploration Satellite Service
ESA	European Space Agency
ITU	International Telecommunication Union
LFM	Linear Frequency Modulated
MMSE	Minimum Mean-Square Error
NESZ	Noise-Equivalent Sigma Zero
NMSE	Normalised Mean-Square Error
Pol-InSAR	Polarimetric SAR Interferometry
RCS	Radar Cross Section
REP	Residual Excess Power
RFI	Radio Frequency Interference
SAR	Synthetic Aperture Radar
SFR	Spectral-Flatness Ratio
SIR	Signal-to-Interference Ratio
SLC	Single Look Complex
SNR	Signal-to-Noise Ratio
TDBP	Time-Domain Back-Projection
TSNB	Time-Stationary Narrowband (interference)
TVWB	Time-Varying Wideband (interference)

Acknowledgements

