SCUOLA DI INGEGNERIA INDUSTRIALE
E DELL'INFORMAZIONE

Executive Summary of the Thesis

RFI Detection and Suppression in P-Band Spaceborne SAR Data

Laurea Magistrale in Space Engineering

Author: Alessandro Bellezza

Advisor: Prof. Stefano Tebaldini

Academic year: 2025-26

1. Introduction

Above-ground biomass is one of the largest and least constrained terms in the global carbon cycle, and the European Space Agency's BIOMASS mission was conceived to map it from space, carrying the first P-band synthetic aperture radar (SAR) ever to operate in orbit. The long P-band wavelength is what reaches the woody structures that hold most of that biomass, but it confines the radar to a narrow 6 MHz allocation around 435 MHz, shared on a secondary basis with a dense population of terrestrial transmitters. Any spaceborne receiver in this band therefore collects energy it was never meant to see, and the mission must treat in-band radio frequency interference (RFI) as a permanent feature of its data rather than an occasional fault. [2]

This thesis develops, implements and validates a data-driven, catalogue-free RFI *detection-and-suppression* chain for BIOMASS commissioning-phase data, built throughout on a single physical fact: the polarimetric structure that the interference corrupts is also what makes it visible and removable. The work yields two original contributions, a coherence-based detector and a scene-preserving coherent-cancellation removal, presented here in the order in which they were built.

2. The Interference Problem at P-Band

The mission's sensitivity requirement (NESZ ≤ -27 dB across the swath) is stringent, and pre-launch airborne campaigns established that unmitigated RFI can degrade the effective sensitivity by 5 to 20 dB in affected areas, on its own enough to push biomass-retrieval errors past the scientific budget. Narrow-band carriers leave bright azimuth streaks at fixed range positions; wide-band bursts raise the noise floor across a whole azimuth segment; both bias the calibrated backscatter and, most critically for the interferometric and tomographic phases of the mission, erode the coherence on which height and vertical-structure estimates depend. Because the cross-polarized channels carry weaker backscatter than the co-polarized ones, a fixed interference power degrades them more severely, precisely the channels that polarimetric biomass retrieval most needs. The scarcity of documented distribution maps for P-band interferers makes a data-driven mitigation, relying on no pre-compiled catalogue, a necessity rather than a preference.

Two facts from radar fundamentals frame the solution. First, transforming each received pulse to the range-frequency / azimuth-time domain (f, τ) separates the two signal classes geometrically: a persistent narrow-band tone (TSNB) appears as a horizontal stripe at a fixed frequency, a wide-band burst (TVWB) as a vertical stripe confined to a few azimuth lines, while the focused scene spreads broadly across the plane. Second, and central to everything that follows, the four focused channels are not acquired independently: the two channels that share a receive antenna port (HH with VH, and VV with HV) collect any external wavefront of that polarization, because it enters through the same hardware. Over flat natural terrain the co- and cross-polarized backscatter are statistically independent, so the cross-channel coherence of a receive pair sits near zero; wherever a shared-receive interferer dominates, that coherence is driven toward unity. This contrast is the seed of the whole work: near zero on clean scene and near one on interference, it is independent of the interferer's waveform and of its absolute power.

State of the art. The established single-channel suppression baseline, formalized by Meyer et al. [1], detects RFI by comparing spectral power against a background model estimated from the data and then notches the flagged samples; it is effective but inherits the weakness of every background-based test, losing calibration once a large fraction of bins is contaminated, and it treats each channel in isolation, penalizing the weak cross-pol channels. The observation that cross-channel coherence could replace the background model was first reported by Uozumi et al. [4], though only as an indicator rather than a complete detector. On the suppression side, coherent cancellation, which uses one channel as a reference to subtract the interference from another in the spirit of the classical Wiener canceller, has been applied to polarimetric SAR by Ding and Wang [5], but with a single global weight that cannot follow interference whose coupling varies across the scene. Each of these external results marks the starting point that one of this thesis' contributions then extends.

3. A Coherence-Based, Catalogue-Free Detector

The first contribution converts the coherence contrast into a complete detection algorithm. Its statistic rests on a single clean relation: when scene, interference and noise are mutually uncorrelated, the cross-channel coherence of a receive pair depends only on the correlation between the interferers in the two channels. Writing each channel as a scene term plus an interference term, $x_1 = s_1 + n_1$ and $x_2 = s_2 + n_2$ (neglecting other noise sources), every contribution to the cross-spectrum cancels except the interference cross-term, and the coherence takes the form

$$|\gamma| = \frac{|\langle n_1\, n_2^* \rangle|}{\sqrt{P_1\, P_2}}, \tag{1}$$

where $\langle n_1\, n_2^* \rangle$ is the cross-correlation of the interference in the two channels and P_1, P_2 are the total channel powers. The coherence therefore stays near zero on clean scene, where no interference is shared, and rises as a correlated interferer comes to dominate the two channels.

Two properties make this an unusually convenient detection signal. It is *waveform-agnostic*, since nothing in the relation depends on the interferer's modulation, bandwidth or duty cycle. And it is *power-normalised*: the coherence is a bounded ratio whose clean baseline does not depend on the local backscatter, so a single threshold applies identically over bright and dark terrain, with no background to estimate from possibly contaminated data. This is why the detector keeps its sensitivity on the weak cross-pol channels, where amplitude statistics become unreliable, and why that sensitivity is in fact naturally higher there, a given interference power being a larger fraction of the weaker channel's total.

The working principle rests on four hypotheses, each tied to a design choice and a failure mode: the two receive-pair channels must carry independent scene backscatter (which holds over distributed natural terrain and is further protected by operating on defocused data, where residual scene correlation is diluted across the plane); the interference must couple identically into both channels (guaranteed by the shared receive port); the one-pulse acquisition delay between the two interleaved channels must be short compared to the interference decorrelation time; and the interference must locally dominate the scene for the coherence to rise. The third hypothesis is the one that later limits the cancellation depth, and it returns as the leading line of future work. Around this statistic the pipeline is built. The focused product is defocused to the (f, τ) domain, where the RFI

is most compact, and the coherence of Eq. (1) is estimated twice, with two oriented sliding windows matched to the two morphologies: an azimuth-elongated window that accumulates a persistent tone coherently while averaging the independent scene toward zero, and a frequency-elongated window that does the same for a wide-band burst; subsequently, each map is binarized through thresholding. The threshold is not heuristic but follows from the estimation bias of the zero-expectation coherence over a finite averaging window, which links window size and threshold directly to the per-pixel false-alarm probability. A line-detection stage then keeps only connected structures longer than a minimum length, discarding isolated false alarms at negligible cost in missed detections, since a genuine tone persists over many pulses and a genuine burst over many frequency bins. The two oriented masks are OR-combined into a single binary mask, shared by both channels of the receive pair to keep the polarimetric correction consistent, and applied in the baseline configuration as a spectral notch.
Set against the single-channel reference, both methods operate in the same domain, target the same two classes with two sub-detectors, and reject false alarms before notching; the difference is concentrated in the decision statistic. Replacing a background-power test with a coherence test removes the dependence on a background estimate entirely, so heavy contamination (which biases the single-channel test) merely enlarges the region where the coherence is close to one. The shared mask lets the stronger channel of a pair support the weaker one, at the price of making a false alarm more likely, therefore the need for a finely tuned rejection stage. Lastly, the pipeline reduces to separable convolutions and connected-component analysis, scaling linearly with the number of pixels.

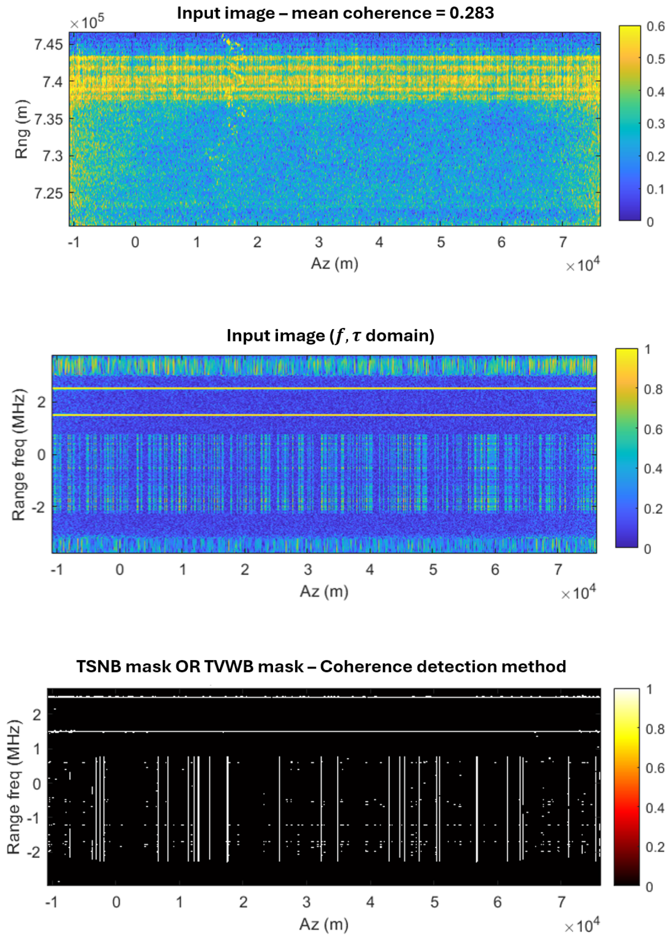

Figure 1: The detector operates in the range-frequency / azimuth-time domain (middle image), where the two interference classes appear as oriented high-coherence structures that the two matched estimators capture.

4. Validation in a Controlled Environment

A detector can be validated only where the truth is known, and on real interference there is neither a clean reference nor exact knowledge of which cells are contaminated. The detector was therefore first applied in a controlled simulation. A clean quad-polarimetric scene, known to be interference-free, is defocused to the raw-equivalent domain, injected with synthetic interference of known type and power exactly as physical interference would enter the receiver, and refocused into the corrupted images fed to every method under test. The reliability of the procedure rests on the invertibility of the defocus–refocus operator pair: the residual of an empty round trip fixes the noise floor against which all reconstruction results are read. The injector reproduces both interference classes with their full parametric variety (tone frequency and range support, burst bandwidth and waveform, azimuth activity) and injects the identical waveform into the two channels that share the same receive polarity, so that the experiments probe the remaining hypotheses

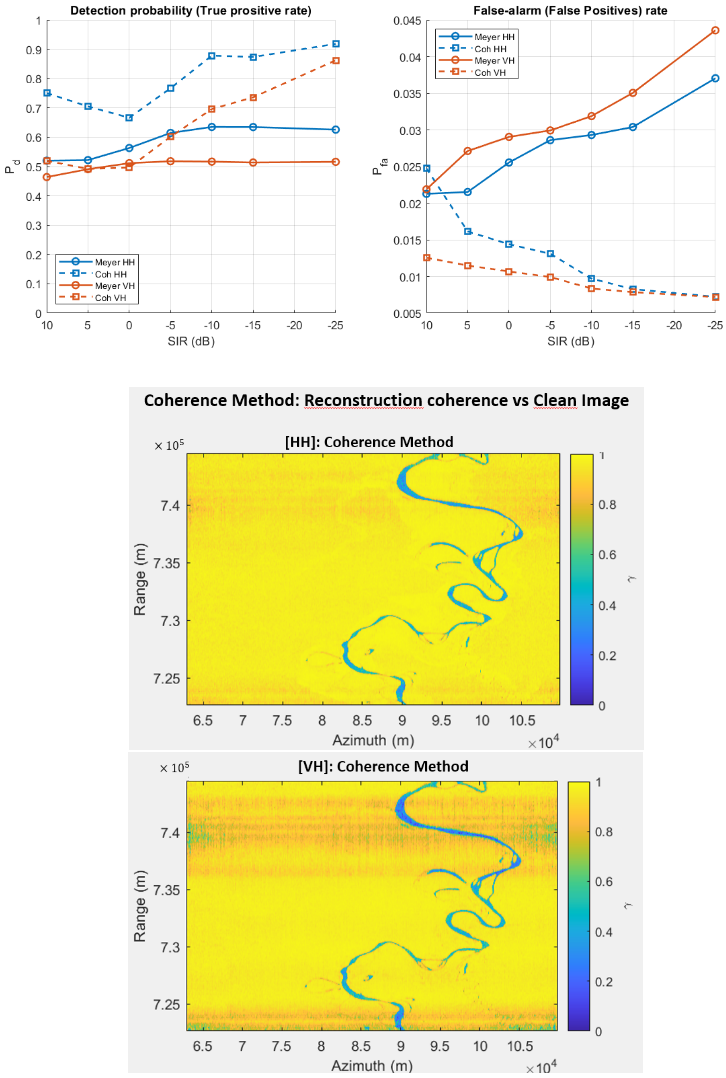

Figure 2: Controlled-simulation validation: SIR sweep results and Reconstruction coherence

and the pipeline rather than a built-in assumption. Three families of metrics follow: detection scores against the known contamination geometry, per-channel reconstruction against the clean reference, and recovery of the polarimetric pair statistics that feed the interferometric and tomographic products. [3]

The controlled campaign delivered a clear verdict. On the mixed scenario (most representative of the real data environment) the coherence detector flagged the wide-band contamination more accurately than the single-channel reference while matching it on the narrow-band tones, and reconstructed both channels more faithfully: on the co-pol channel the reconstruction coherence rose from 0.92 of the reference method to 0.94 with coherence-based detection, and the surviving interference power fell from 19.4% to 10% of the input. The margin was even wider on the cross-pol channel, where the mean coherence rose from 0.81 to 0.91. The same dominance appeared in detection probability across the full severity sweep, the gap widening on the cross-pol channel where the single-channel test reaches a plateau. The experiment thus established the result the rest of the thesis builds on: the coherence detector is at least comparable to the reference even under conditions favourable to the latter, and decisively better in the two regimes that matter for BIOMASS, the weak cross-pol channels and the heavily contaminated acquisitions, with no background model and no prior knowledge of the interferer. It also exposed an intrinsic ceiling: under notch removal every correctly flagged cell is a cell whose scene content is destroyed, so beyond a certain contamination extent no detect-and-notch scheme, however accurate its mask, can recover the channel. Overcoming that ceiling required a different removal operator.

5. Suppression on Real BIOMASS Data

The second contribution, a coherent subtraction-based removal, counteracts the notch limitation and is validated on real BIOMASS acquisitions independent of those used to tune the detector. It reuses the same physical fact that powers the detection. Taken column by column in the (f, τ) domain, the two co-located channels of a receive pair are written as a useful scene term plus an interference term, $x_1 = s_1 + n_1$ for the co-pol (*master*) channel and $x_2 = s_2 + n_2$ for the cross-pol (*auxiliary*) one. The removal forms the filtered estimate $\hat{Y} = X_1(f) + W(f)\, X_2(f)$ and chooses the complex weight $W(f)$ that minimises the coherence between $\hat{Y}$ and the auxiliary x_2, equivalently nulling their cross-spectrum. Imposing:

$$S_{\hat{Y}x_2} = S_{x_1x_2} + W\, S_{x_2x_2} = 0 \qquad (2)$$

From which results:

$$W(f) = -\frac{S_{x_1x_2}(f)}{S_{x_2x_2}(f)}. \qquad (3)$$

This weight isolates the interference rather than the scene under three hypotheses that the polarimetric acquisition geometry makes natural:

(i) the useful scene components of the two channels are mutually uncorrelated, which holds for distributed natural terrain; (ii) the interference components are mutually correlated, because they enter through the same receive port; and (iii) the auxiliary channel is dominated by the interference, a condition the cross-pol channel satisfies through its lower signal-to-interference ratio. Under (i) to (iii) the cross-spectrum collapses onto the interference cross-term and the weighted reference converges to the interference present in the master channel, so that subtracting WX_2 cancels N_1 while leaving the independent scene term s_1 in place, exactly the cell content a spectral notch would discard. The cross-pol channel, having no third independent reference, is still notched.

In the focused image each range-frequency column is the Fourier transform of one received pulse, so the same reasoning applies per column, the weight estimated as a local slow-time average,

$$W(f,x) = -\frac{\langle \mathrm{HH}(f,x)\,\mathrm{VH}^*(f,x)\rangle}{\langle |\mathrm{VH}(f,x)|^2\rangle} \qquad (4)$$

over a short azimuth-time window. Two safeguards keep the estimate trustworthy and distinguish it from the global-weight cancellation of the prior art: it is gated by the validated detection mask, so it never acts where the auxiliary is scene-dominated and hypothesis (iii) fails; and it is estimated locally, tracking the spatial and spectral variation of the coupling that a single global weight cannot follow. When the auxiliary is not fully interference-dominated, at high SIR, the residual scene power retained in the denominator of Eq. (4) causes a slight under-correction and signal distortion, which restricting the average to small windows around the detected high-coherence cells keeps negligible.

A benchmark was run against four alternatives on a strongly contaminated, mixed-class scene: the single-channel Meyer detector with notching, the coherence detector with notching, a unit-weight variant to test the necessity of a local weight estimation, the proposed local-Wiener cancellation, and, as an external end-to-end comparison, the official ESA BIOMASS filter currently deployed. With no clean reference available, the comparison rests on indicators read across methods: the residual cross-channel coherence at the flagged cells, a set of spectral-structure indices, and visual inspection of focused images and spectra.

Table 1: Residual cross-channel coherence at the flagged cells, real data (lower is better).

Method	HH–VH	VV–HV
Input	0.52	0.53
Meyer + notch	0.476	0.478
Coherence + notch	0.416	0.429
Local Wiener cancellation	**0.34**	**0.34**
Operational ESA filter	0.48	0.51

The results separate the two stages of the chain. Holding the removal fixed at notching, swapping the single-channel detector for the coherence detector already lowers the residual coherence at the flagged cells (from 0.476 to 0.416 on the H-receive pair, from 0.478 to 0.429 on the V-receive pair) and flags the interference far more completely, about a 20% increase in mask coverage, using one shared mask per pair rather than two independent ones.

Holding the detector fixed and varying the removal then isolates the cancellation's effect. Subtraction-based removal reaches the lowest residual coherence of all (around 0.34), well below the notch, because actively cancelling the common reference decorrelates the pair more effectively than zeroing the cells. Overall, the local Wiener weight achieves the best combined result, keeping the low residual coherence while bringing the spectral structure back below the input and leaving the focused image free of the original interference stripes; its only cost is runtime, being the slowest method on this scene.

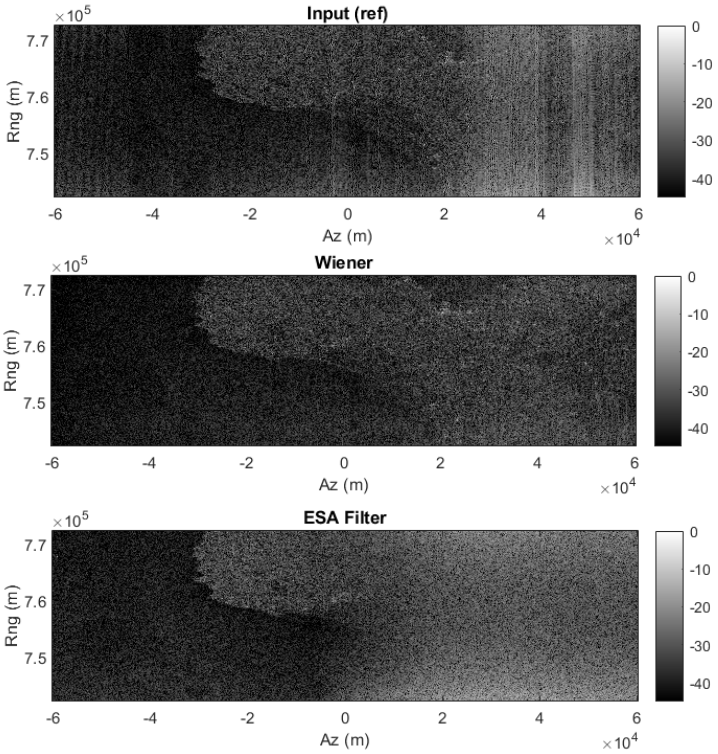

Figure 3: Real-data comparison on a strongly contaminated coastal scene (HH Channel Amplitude).

Against the operational ESA filter the proposed method results promising. The ESA output is visually very smooth and comparable to the Wiener result on cleanliness, but it achieves this by whitening the range-power spectrum to a near-constant level and smoothing the time-domain image, and it reduces the cross-channel coherence at the flagged cells only marginally, by a few percent. The Wiener cancellation, acting only where the interferer is actually present, suppresses the spectral peaks and the associated visual artifacts and roughly halves the residual coherence; in the visual comparison the cleaned co-pol channel is the cleanest of all while preserving the most scene detail, ESA included. Read as a map of the local P-band environment, the detector output also characterises the scene's RFI directly, resolving a swept wide-band emitter and a persistent narrow-band population overlaid with both range frequency- and azimuth-localized bursts.

6. Conclusions

This thesis shows that the polarimetric structure of the BIOMASS data is not merely a product to be protected from interference but the means to detect and characterize it: the same shared-receive coupling that makes the cross-pol channels vulnerable to RFI makes one channel an effective reference for cleaning another. Building both the detection and the suppression on this single physical fact yields a chain that needs no prior knowledge of the interferers, stays effective on the weak channels and under heavy contamination where amplitude-based methods are disadvantaged, and preserves the scene content that notching discards. BIOMASS being the first spaceborne P-band SAR, the developed method is among the first methods reported in the open literature to detect and suppress RFI directly on real spaceborne P-band data, alongside the mission's own operational baseline. Two directions of continuation for this work follow naturally from the analysis: compensating the one-pulse inter-channel delay, that currently caps the cancellation depth, is expected to bring the Wiener residual to the notch level while keeping its scene-preservation advantage. Secondly, exploiting the proposed method lower computational cost at high pixel number could highlight its predicted efficiency, especially when applied to larger datasets that will be available in the future.

References

[1] Franz J. Meyer, Jeremy B. Nicoll, and Anthony 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.

[2] Shaun Quegan, Thuy Le Toan, Jérôme Chave, Jørgen Dall, Jean-François Exbrayat, Dinh Ho Tong Minh, Mark Lomas, Mauro Mariotti D'Alessandro, Philippe Paillou, Konstantinos Papathanassiou, Fabio Rocca, Sassan Saatchi, Klaus Scipal, Hank Shugart, T. Luke Smallman, Maciej J. Soja, Stefano Tebaldini, Lars Ulander, Ludovic Villard, and Mathew Williams. The European Space Agency BIOMASS mission: Measuring forest above-ground biomass from space. *Remote Sensing of Environment*, 227:44–60, 2019.

[3] 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, January 1999.

[4] Akira Uozumi, Akira Hirose, and Ryo 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.

[5] Junxu Wang, Lei Zhang, Zhichao Meng, Shaopeng Wei, and Li-Li Chen. General RFI suppression with sidelobe cancellation filtering for dual-polarization SAR images. *IEEE Geoscience and Remote Sensing Letters*, 20:4012005, 2023.