Doctorate of National Interest

in

Space Science and Technology

Advanced Radar Techniques for Planetary Exploration: 3D Subsurface Imaging of Mars and Near-Earth Asteroid Characterization

Alessio Margheri

Istituto Nazionale di Astrofisica (INAF)

Co-Supervisors
Giuseppe Pupillo
Istituto Nazionale di Astrofisica
Lorenzo Bruzzone
University of Trento

Supervisor
Roberto Orosei
Istituto Nazionale di Astrofisica

Academic Year 2025–2026

March 2026

Abstract

Planetary radar sounding enables the investigation of subsurface structures and three-dimensional characterization of celestial bodies that are otherwise inaccessible through optical observations alone. This thesis addresses advanced radar signal processing techniques for two complementary applications: subsurface imaging of Mars through MARSIS (Mars Advanced Radar for Subsurface and Ionosphere Sounding) data, and three-dimensional shape reconstruction of near-Earth asteroids using ground-based radar observations.

The theoretical foundation is established through a comprehensive treatment of electromagnetic wave propagation, surface scattering mechanisms, and radar equation fundamentals. Particular emphasis is placed on digital signal processing techniques for radar sounding, including matched filtering, pulse compression, and synthetic aperture radar (SAR) principles. The framework extends to advanced focusing algorithms, with detailed analysis of both time-domain and frequency-domain backprojection methods, inverse scattering theory, and tomographic reconstruction approaches. This theoretical synthesis is necessary given the vast and fragmented nature of radar imaging theory across different application domains.

The investigation of Mars subsurface focuses on MARSIS data from the south polar region, specifically the Ultimi Scopuli area where previous studies identified anomalously bright basal reflections potentially consistent with subsurface liquid water at an apparent depth of 3 km. Through comparison with simulated data, a remnant phase term was discovered causing incoherent integration of orbital data. This led to a complete revision of the raw MARSIS processing chain, identifying noise sources including demodulation artifacts, orbital radial velocity variations, and ionospheric phase delays. The core contribution consists in the development and implementation of a three-dimensional frequency-domain backprojection algorithm, specifically optimized for MARSIS acquisition geometry and data characteristics. Applied first to a subset of 3 orbits selected for their positive focusing gain, the algorithm produces a volumetric reconstruction of the subsurface reflectivity that is qualitatively improved with respect to the uncorrected case, with spatially separated bright spots consolidating into a single, geometrically connected reflector after phase and noise correction. The reconstruction pipeline was subsequently extended to the full set of 27 noise- and phase-corrected orbits over the Ultimi Scopuli region. The resulting 3D volume reveals a coherent, elongated basal reflector at an apparent depth of 3 km whose geometry and spatial extent are qualitatively consistent with the map of anomalously high basal dielectric permittivity reported by [1], providing qualitatively and independent three-dimensional geometric evidence for the spatial continuity of the candidate subsurface structure.

Ground-based radar is a powerful tool for the characterization of Near-Earth Objects (NEOs), using both continuous-wave (CW) and delay-Doppler imaging techniques. Contributing to the European Space Agency pilot project "NEO Observation Concepts for Radar Systems" and to the subsequent activities, radar observations have been conducted using several radio telescopes in

Europe as receivers, including the 32m "G. Grueff" Medicina antenna. For asteroid 2005 LW3, high-resolution integrated spectra revealed a binary system through the detection of a secondary spectral peak. In this work, a progressive implementation of three-dimensional shape reconstruction algorithms was developed to address scenarios where no range information is available, evolving from two-dimensional convex hull models through rotational ellipsoid fitting, radial and tangential perturbation models, culminating in a sophisticated adaptive vertex model capable of locally augmenting vertex density by optimizing residuals against spectral data. The reconstruction demonstrates that accurate 3D characterization is achievable using only Doppler data from CW observations.

Additionally, this work includes pioneering efforts in Very Long Baseline Interferometry (VLBI) for asteroid radar observations. A complete feasibility study was performed for the observation of asteroid 2025 FA22, including SNR estimates for each telescope and baseline, as well as simulations of observational visibility conditions. This represents one of the first VLBI radar observations of an asteroid, with successful detection confirmed by Lovell telescope data, establishing a foundation for future interferometric radar campaigns.

This research advances planetary radar sounding through the development of robust signal-processing algorithms tailored to challenging acquisition scenarios for subsurface exploration on Mars, while it also implements innovative radar-based 3D modeling algorithms for near-Earth objects, supporting both scientific investigation and planetary defense applications.

Keywords: radar, MARSIS, NEO, asteroid, Mars

Contents

Chapter 1

Basic radar principles

1.1 What is a radar?

Radar (RAdio Detection And Ranging) is an electrical system that transmits radiofrequency (RF) electromagnetic (EM) waves toward a target of interest and receives and detects the reflected EM waves. Figure **1.1** shows the principal subsystems that are part of a radar system, although each system can be different according to the purpose it has been designed for:

- the transmitter generates the EM waves;
- the antenna takes as input the EM waves generated by the transmitter and introduces them into the propagation medium;
- the receiver collects the reflected signal from the target;
- the Signal Processor is the device that sorts and analyzes the data collected by the receiver.

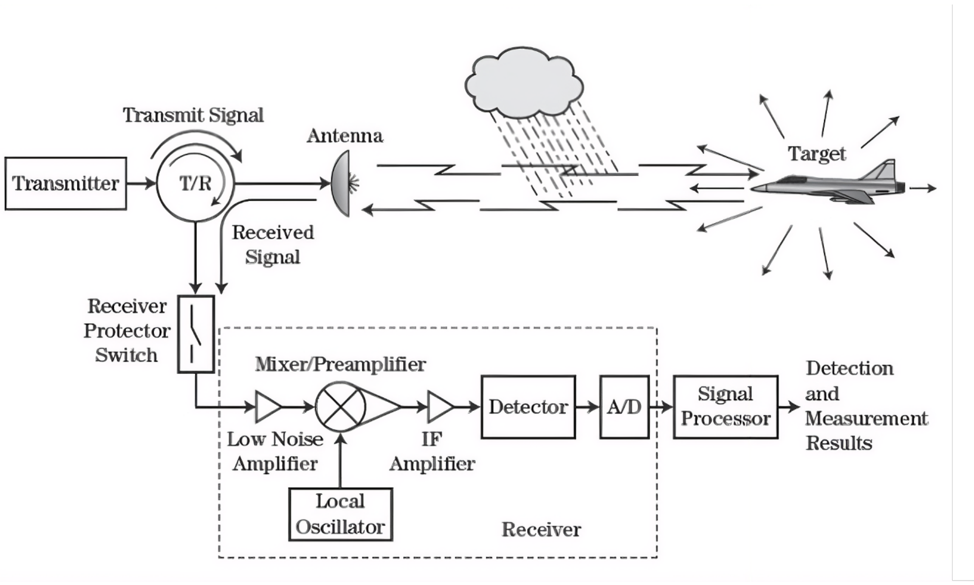

Figure 1.1: Major elements of the radar transmission/reception process. From [2].

The transmit/receive (T/R) device (also called *circulator*) connects the transmitter to the antenna. The T/R allows to connect both the transmitter and the receiver to the antenna simultaneously and at the same time it isolates the receiver from the transmitter to protect the sensitive receiver components from the high-powered transmitted signal.

The physics of a radar observation is straightforward. The transmitted signal propagates through the medium and reaches the target. The impinging EM waves generates currents on the target, that are reradiates back into the environments. The backscattered signal acquired by the receiver contains the desired signal reflected by the target and other unwanted components, such as echoes from other surfaces on the ground illuminated by the transmitted signal and from the propagation medium (such as the ionosphere). These components are called *clutter*, and they may alter the strength and the shape of the reflected signal.

Once the radio frequency (RF) signal is collected by the receiver, it is amplified, converted to intermediate frequency (IF), digitized by an analog-to-digital converter (ADC), and subsequently passed to the signal/data processor. The detector removes the carrier from the modulated target return signal before it is given as input to the Signal Processor. This device is responsible for determining the presence of a target in the presence of interference, which can be in four different forms: internal and external noise, clutter, electromagnetic interferences (EMI), which are unintentional external EM waves created by other human-made sources, and intentional jamming from an electronic countermeasures (ECM) system, in the form of noise or false targets.

1.2 Electromagnetic background

1.2.1 Waves equation

To understand the principle under which radar observations work, it is necessary to recall some concepts from EM theory. In particular, it is important to understand how EM waves interact with material media. The interaction may involve scattering, absorption, transmission, and emission, or combinations thereof. All these phenomena will be briefly analyzed in this Chapter.

Electromagnetic waves are electric and magnetic field waves, oscillating at the carrier frequency. The nature of the EM fields is the described by Maxwell's equation. For time-harmonic fields in a source-free, linear, homogeneous, and isotropic medium, Maxwell's equations in the frequency domain are:

$$\nabla \cdot \mathbf{E} = 0 \tag{1.1}$$

$$\nabla \times \mathbf{E} = -j\omega\mu\mathbf{H} \tag{1.2}$$

$$\nabla \cdot \mathbf{H} = 0 \tag{1.3}$$

$$\nabla \times \mathbf{H} = j\omega\varepsilon_r\varepsilon_0\mathbf{E} \tag{1.4}$$

where $\mathbf{E}$ and $\mathbf{H}$ are the electric and magnetic field vectors, respectively; ω is the angular frequency; $\mu = \mu_r\mu_0$, with $\mu_0 = 4\pi \times 10^{-7}$ H/m being the vacuum permeability and μ_r the relative permeability of the medium; ε is the relative electric permittivity (the terms "permittivity" and "dielectric constant" will be used interchangeably throughout this work to refer to relative permittivity); $\varepsilon_0 = 8.8541878188 \cdot 10^{-12}\mathrm{Fm}^{-1}$ is the vacuum permittivity; j denotes the imaginary unit; and ∇ is the vector differential operator. Usually the relative permittivity ϵ takes a value in the range between 1 to 80 for most geological materials , while $\mu = 1$ for non magnetic geologic materials, like Mars surface and subsurface.

If we take the curl of both sides of Equation **1.2** and then use appropriate substitution it is possible to write the homogeneous wave equation for $\mathbf{E}$, namely:

$$\nabla\mathbf{E} + \omega^2\mu_0\mu_r\epsilon\epsilon_0\mathbf{E} = 0 \tag{1.5}$$

By defining the propagation constant γ as:

$$\gamma^2 = -\omega^2 \mu \epsilon \epsilon_0 \tag{1.6}$$

Equation **1.5** can be written as:

$$\nabla^2 \mathbf{E} - \gamma^2 \mathbf{E} = 0 \tag{1.7}$$

Vice versa, taking the curl of both sides of Equation **1.4** and then use Equation **1.2** to eliminate **E**, one obtains a wave equation for **H**:

$$\nabla^2 \mathbf{H} - \gamma^2 \mathbf{H} = 0 \tag{1.8}$$

The resulting wave equations demonstrate that, under the assumptions of linearity, homogeneity, and absence of sources, electromagnetic fields propagate as waves governed by the intrinsic properties of the medium. In particular, the propagation constant γ encapsulates the combined effect of angular frequency and electromagnetic parameters, providing a compact description of the propagation characteristics relevant to the interpretation of radar signals. In this context, electromagnetic waves are commonly approximated as plane waves. This approximation is valid when the waves propagate in a homogeneous medium and the observation point is sufficiently far from the source, such that the curvature of the wavefronts can be neglected over the region of interest. The far field condition is thus expressed as:

$$d > \frac{2D_{\text{max}}^2}{\lambda} \tag{1.9}$$

where d is the distance from the source, D_{max} is the antenna dimension, and λ is the observation wavelength.

Under these conditions, the electric and magnetic fields are transverse to the direction of propagation and can be described by simple traveling-wave solutions of the homogeneous wave equations. Plane waves are good approximations to real waves in many practical situations, considering also that more complicated EM wave fronts can be treated as a superimposition of plane waves.

A solution to the first wave equation of particular interest to radar is a traveling sinusoidal electric field wave, with the amplitude of the z directional component having the form:

$$E = E_0 \cos\left(kz - \omega t + \phi\right) \tag{1.10}$$

where $k = 2\pi/\lambda$ is the wavenumber, λ is the wavelength in the medium and ϕ is the phase of the wave. k may be considered a constant of the medium for a particular frequency and may also be referred to as the propagation factor for a medium. It can be shown that in free space, **E** and **H** of the propagating EM wave are orthogonal and the direction of the propagating wave is orthogonal to both **E** and **H**. The direction of propagation is given by the time-averaged Poynting vector:

$$\mathbf{S} = \frac{1}{2}\text{Re}\{\mathbf{E} \times \mathbf{H}^*\} \tag{1.11}$$

whose magnitude represents the intensity of the wave. The ratio $|\mathbf{E}|/|\mathbf{H}|$ defines the wave impedance of the medium:

$$\eta = \sqrt{\frac{\mu}{\varepsilon}} \tag{1.12}$$

which in free space takes the value $\eta_0 = \sqrt{\mu_0/\varepsilon_0} \approx 377\,\Omega$. Since **E** and **H** are related through η, the EM wave is typically described in terms of the electric field alone.

1.2.2 Polarization

The polarization of a uniform plane wave describes the locus traced by the tip of the **E** vector (in the plane orthogonal to the direction of propagation) at a given point in space as a function of time [3]. Let's consider an EM wave traveling in the $+z$ direction, as shown in Figure **1.2**. As stated before, the z components of the electric and magnetic fields of a z-propagating plane wave are both zero, leading to:

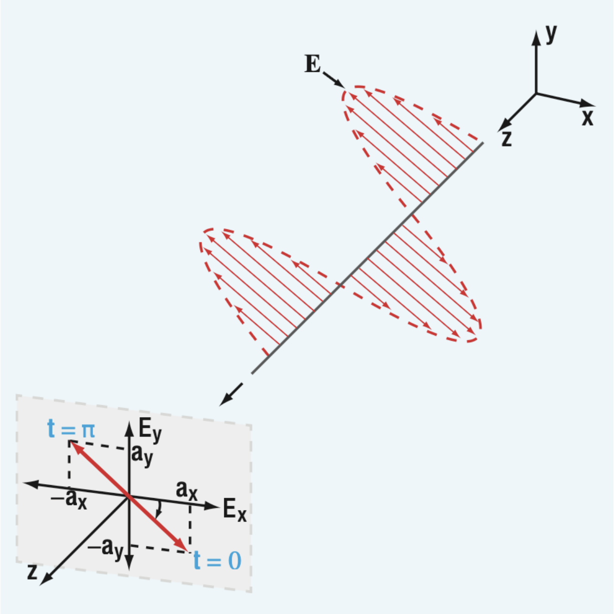

Figure 1.2: Linearly polarized wave traveling in the $+z$ direction. From [4].

$$\mathbf{E}(z) = (\hat{\mathbf{x}}E_{x0}(z) + \hat{\mathbf{y}}E_{y0}(z))e^{-jkz}e^{j\omega t} \tag{1.13}$$

The two amplitudes E_{x0} and E_{y0} are complex quantities, characterized by a magnitude and a phase angle. The exponential term function of space describes the effect on phase and amplitude (i.e., attenuation) of the signal as it propagates through the medium. Let's assign for convenience E_{x0} a phase of zero and a phase δ to E_{y0}, relative to that of E_{x0}:

$$E_{x0} = A_x, \tag{1.14}$$

$$E_{y0} = A_y e^{j\delta} \tag{1.15}$$

The amplitude of the electric field can be written as:

$$\mathbf{E}(z,t) = \mathrm{Re}[\mathbf{E}(z)e^{i\omega t}] = \hat{\mathbf{x}}A_x \cos(\omega t - kz) + \hat{\mathbf{y}}A_y \cos(\omega t - kz + \delta) \tag{1.16}$$

and its orientation in space is given by the orientation angle ϕ_p:

$$\phi_p(z,t) = \tan^{-1}\left[\frac{A_y \cos(\omega t - kz + \delta)}{A_x \cos(\omega t - kz)}\right] \tag{1.17}$$

and depends on the relative amplitude and phase of the x and y components in case of an isotropic medium. i.e., the wavenumber of the medium is the same for any orientation of the incoming electric field.

A wave is said to be *linearly polarized* if $E_x(z,t)$ and $E_y(z,t)$ are in phase ($\delta = 0$) or out of phase ($\delta = \pi$). Under these assumption, the orientation angle becomes, for the out of phase case:

$$\phi_p(z,t) = \tan^{-1}\left[\frac{-A_y \cos(\omega t - kz)}{A_x \cos(\omega t - kz)}\right] = \tan^{-1}\left(-\frac{A_y}{A_x}\right) \tag{1.18}$$

The orientation angle becomes a constant value, independent of both z and t. This means that the electric field always lie within a plane of constant orientation in space. The same holds for the in phase case. When the phase difference is $\delta = \pm\pi/2$ and the magnitudes of the x and y components of $\mathbf{E}$ are equal, the signal is circularly polarized. The wave polarization is called *left-hand circular* (LCP) if $\delta = \pi/2$ and *right-hand circular* (RCP) if $\delta = -\pi/2$. The sign of the phase shift indicates whether the electric vector rotates clockwise or counter-clockwise. In this scenario, the orientation angle becomes:

$$\phi_p(z,t) = \tan^{-1}\left[\frac{\pm A\cos(\omega t - kz)}{A\cos(\omega t - kz)}\right] = \pm(\omega t - kz) \tag{1.19}$$

Thus, the orientation angle of $\mathbf{E}$ is dependent on both space and time, rotating at angular frequency ω. In case in which $|A_x| \neq |A_y|$, the signal is elliptical polarized.

Usually in remote sensing applications radar instruments employ a linear or a circular polarized signal. MARSIS, a nadir looking radar, for instance, utilizes a linear polarized wave, while in bistatic observation it is usually preferred a circular polarized signal, to take advantage of the oblique geometry, like Medicina Radiotelescope does.

1.3 Propagation in a medium

As stated before, Maxwell's Equations (**1.1**), (**1.2**), (**1.3**), (**1.4**) describe the propagation of EM waves. In free space the magnetic permeability and electric permittivity are constants, meaning that they are independent of frequency and the medium is not dispersive. On the contrary, EM waves propagating through media experience losses, to both the electric and magnetic field, which causes attenuation of the original signal.

In remote sensing observation, measuring the dielectric constant of a material is crucial, since it depends on its composition, physical condition (e.g. temperature, pressure, density, and phase) as well as the wavelength of the observation. In general the complex permittivity ϵ may be expressed as:

$$\epsilon = \epsilon' - j\epsilon'' = \epsilon' - j\frac{\sigma}{\omega\varepsilon_0} \tag{1.20}$$

where ϵ' and ϵ'' are the real and imaginary parts of the complex permittivity, respectively, σ is the material conductivity, and $\omega = 2\pi f$ is the angular frequency. In particular, the real component ϵ' varies with the wavelength of the EM wave and quantifies the material's ability to store electrical energy through polarization mechanisms. The imaginary component ϵ'' accounts for energy dissipation via dielectric relaxation and conduction losses. Media with $\epsilon'' \gg \epsilon'$ are said to be good conductors, while materials with $\epsilon'' \ll \epsilon'$ are called dielectric. As shown in the formula, ϵ'' depends from the instrument's wavelength and the material's conductivity. Consequently, in radio frequency remote sensing of planetary surfaces, the imaginary component of the dielectric constant demonstrates stronger wavelength-dependent variations than its real counterpart [5]. Also, we recall that a material is defined as lossless if $\sigma = 0$. The loss tangent δ provides a measure of the material's overall dissipative behavior:

$$\tan\delta = \frac{\epsilon''}{\epsilon'} \tag{1.21}$$

Now let's recall the definition of the propagation constant γ. Since it's a complex quantity, it may be expressed as:

$$\gamma^2 = -\omega^2 \mu \epsilon_0 (\epsilon' - j\epsilon'') = -\omega^2 \mu \epsilon_0 \epsilon' \left(1 - \frac{\epsilon''}{\epsilon'}\right) \tag{1.22}$$

$$\gamma = \alpha + j\beta \tag{1.23}$$

where α is the medium's attenuation constant and β its phase constant:

$$\alpha = -\omega\sqrt{\mu\epsilon_0}\ \mathrm{Im}\ \sqrt{\epsilon} \tag{1.24}$$

$$\beta = \omega\sqrt{\mu\epsilon_0}\ \mathrm{Re}\ \sqrt{\epsilon} \tag{1.25}$$

Consider the right hand form of Equation **1.22**. For $|x| \ll 1$, the function $(1-x)^{1/2}$ can be approximated by the first two terms of its binomial series; that is, $(1-x)^{1/2} \simeq 1 - x/2$. By applying this approximation for a dielectric medium with $x = j\epsilon''/\epsilon''$ and $\epsilon''/\epsilon' \ll 1$, we obtain:

$$\gamma \approx j\omega\sqrt{\mu_0 \epsilon' \epsilon_0}\left(1 - j\frac{\epsilon''}{\epsilon'}\right) \approx j\frac{2\pi}{\lambda}\sqrt{\epsilon'}\left(1 - j\frac{\epsilon''}{\epsilon'}\right) \tag{1.26}$$

The real and imaginary part of Equation **1.26** are:

$$\alpha \approx \frac{\pi\epsilon''}{\lambda\sqrt{\epsilon'}} \tag{1.27}$$

$$\beta \approx \frac{2\pi}{\lambda}\sqrt{\epsilon'} \tag{1.28}$$

Also, it can be easily shown that $\gamma = \alpha + j\beta = jk$ [6]. Consider now a uniform plane wave with electric field $\mathbf{E} = \hat{\mathbf{x}} E_x(z)$ traveling along the $+z$ direction. From Equation **1.13**, the magnitude of $E_x(z)$ is given by:

$$|E_x(z)| = |E_{x0} e^{-\alpha z} e{-}j\beta z| = |E_{x0}| e^{-\alpha z} \tag{1.29}$$

meaning that the magnitude of an EM wave traveling on the z axis decreases exponentially, with its rate depending on the value of α. Specifically, it defined the *skin depth* δ_s of a medium, as the value of the distance z at which:

$$\frac{|E_x(z)|}{|E_{x0}|} = e^{-1}, \text{with} \tag{1.30}$$

$$\delta_s = \frac{1}{\alpha} \tag{1.31}$$

The skin depth is an indicator of penetration depth for surface-penetrating radar systems. However, the effective penetration depth is influenced by several additional factors, most notably the target's reflection strength and the system's clutter suppression capabilities. In a perfect dielectric, $\sigma = 0$ and $\epsilon'' = 0$, which yields $\alpha = 0$ and therefore $\delta_s = \infty$. Thus, in free space, a plane wave can propagate indefinitely with no loss in magnitude. On the other extreme, in a perfect conductor, $\sigma = \infty$, $\alpha = \infty$ and hence $\delta_s = 0$. Thus, the electric field is confined to the surface of a perfect conductor.

1.3.1 Reflection and Transmission

As shown in Figure **1.3**, an EM wave traveling through a medium that encounters the interface between two media with different dielectric permittivity generates two EM waves: one is transmitted (or refracted) into the second medium, with a different trajectory with respect to the original one;

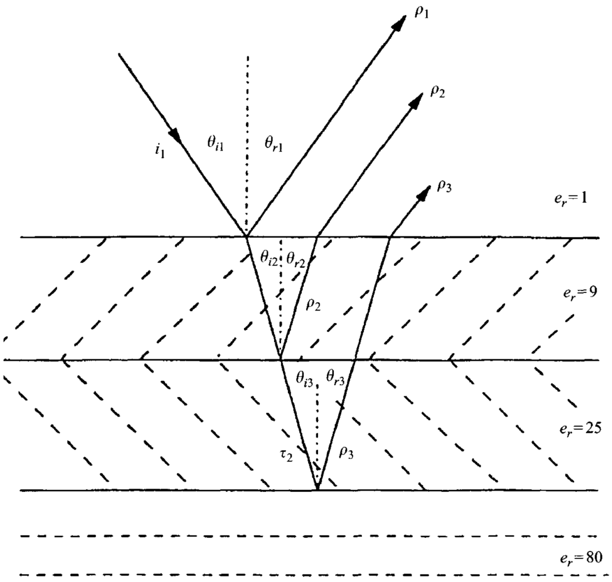

Figure 1.3: EM wave propagation through different media. From [7].

the other is reflected back into the first medium. In the following analysis, we consider a linearly polarized plane wave impinging on a smooth surface with an oblique incidence. The media are assumed to be low-loss. We indicate with θ_{i1}, θ_{r1} and θ_{i2} the angles of incidence, reflection, and transmission (or refraction), defined with respect to the normal to the boundary. The relation between these three angles is defined by Snell's laws of reflection and refraction. The first one states that the angle of reflection equals the angle of incidence, while the second one relates the incidence and transmission angles with the two different phase velocities in the media:

$$\theta_{i1} = \theta_{r1} \tag{1.32}$$

$$n_1 \sin\theta_{i1} = n_2 \sin\theta_{i2} \tag{1.33}$$

where n is the refraction index of a medium, i.e. the ratio between the phase velocity in free space and the phase velocity in the medium. For nonmagnetic materials, $n = \sqrt{\epsilon'}$.

Considering the transmitted wave, if the EM wave interacts with the surface at normal incidence ($\theta_{i1} = 0$) it follows that $\theta_{i2} = 0$. However, at oblique incidence $\theta_{i2} < \theta_{i1}$ when $n_2 > n_1$, and $\theta_{i2} > \theta_{i1}$ when $n_2 < n_1$. A material is said to be more dense than another material if it has a greater refractive index n [3]. Thus, if a wave is incident on a more dense medium, the transmitted wave refracts inwardly, and the opposite is true if a wave is incident on a less dense medium.

In case in which $\theta_{i2} = \pi/2$, there's no transmission in the medium and the refracted wave flows along the surface. The corresponding value θ_c for the incidence angle is called *critical angle*:

$$\sin\theta_c = \frac{n_2}{n_1} \sin\theta_{i2}|_{\theta_{i2}=\pi/2} = \frac{n_2}{n_1} = \sqrt{\frac{\epsilon'_2}{\epsilon'_1}} \tag{1.34}$$

For incidence angles greater than θ_c the incidence wave is totally reflected, while the refracted wave becomes a nonuniform surface wave.

Apart from the directions of propagation it is important to evaluate the amplitudes of the reflected and transmitted waves. To do so, from Maxwell's equations we apply the boundary conditions at the interface between the two media, also called the *phase matching*. Specifically, the tangential component of the total electric field is always continuous across a boundary between two contiguous media, and in the absence of current sources at the boundary, the same is true for the total magnetic field. By applying these assumptions, it is possible to define the *Fresnel reflection and transmission coefficients*:

$$E_0^{r1} = \rho E_0^{i1} \tag{1.35}$$

$$E_0^{i_2} = \tau E_0^{i1} \tag{1.36}$$

The explicit expressions for ρ and τ depends on the angle of incidence and the incident wave polarization. In case of normal incidence, the two coefficients are independent of the polarization, since both the electric and magnetic fields of a normally incident plane wave are tangential to the boundary despite the wave polarization. This is not valid for obliquely incident waves. Defining the *plane of incidence* as the plane containing the normal to the boundary and the direction of propagation of the incident wave, an EM wave with an arbitrary polarization may always be defined as the superposition of two orthogonally polarized waves, one with its electric field parallel to the plane of incidence (parallel polarization) and the other with its electric field perpendicular to the plane of incidence (perpendicular polarization). In the following we provide the formula for τ and ρ for the horizontal polarized case (when $\mathbf{E}$ is perpendicular to the plane of incidence):

$$\rho_{HH} = \frac{\cos\theta_{i1} - \sqrt{(\epsilon_2/\epsilon_1) - \sin^2\theta_{i1}}}{\cos\theta_{i1} + \sqrt{(\epsilon_2/\epsilon_1) - \sin^2\theta_{i1}}} \tag{1.37}$$

$$\tau_{HH} = 1 + \rho_{HH} \tag{1.38}$$

and the vertical polarized case (when $\mathbf{H}$ is perpendicular to the plane of incidence):

$$\rho_{VV} = \frac{(\frac{\epsilon_2}{\epsilon_1} - \sin^2\theta_{i1})^{1/2} - \frac{\epsilon_2}{\epsilon_1}\cos\theta_{i1}}{(\frac{\epsilon_2}{\epsilon_1} - \sin^2\theta_{i1})^{1/2} + \frac{\epsilon_2}{\epsilon_1}\cos\theta_{i1}} \tag{1.39}$$

$$\tau_{VV} = (1 + \rho_v)\frac{\cos\theta_{i1}}{\cos\theta_{i2}} \tag{1.40}$$

The angle at which $\rho_{VV} = 0$ is called *Brewster angle*, defined as the angle at which the vertically polarized component of the incident wave is totally transmitted into the other medium:

$$\theta_B = \sin^{-1}\sqrt{\frac{1}{1 + (\epsilon_1/\epsilon_2)}} = \tan^{-1}\sqrt{\frac{\epsilon_2}{\epsilon_1}} \tag{1.41}$$

it is also known as the polarization angle, since the reflection coefficient for the polarization parallel to the plane of incidence vanishes, while the perpendicular component is fully transmitted. This occurs because the induced electric dipoles radiate no energy in the direction of the reflected wave. As a result, an unpolarized incident wave gives rise to a fully polarized reflected wave.

If we consider an incoming circularly polarized wave, the behavior of the reflected and the transmitted wave is different. the reflection splits into two components: one that remains circularly polarized in the same sense as the incoming wave (referred to as same-sense circularly polarized, SC) and another that is polarized in the opposite or orthogonal sense (referred to as opposite-sense circularly polarized, OC). The transmitted wave largely retains the same circular polarization as the incident wave. Opposite-sense circularly polarized (OC) components in transmission are negligible for isotropic, non-magnetic media [8]:

$$\rho_{SC} = \frac{\rho_{VV} + \rho_{HH}}{2} \tag{1.42}$$

$$\rho_{OC} = \frac{\rho_{VV} - \rho_{HH}}{2} \tag{1.43}$$

To obtain information about the amount of transmitted and reflected power at the interface, it is possible to define the *reflectivity* Γ and the *transmissivity* T as:

$$\Gamma = |\rho|^2 \tag{1.44}$$

$$T = 1 - \Gamma \tag{1.45}$$

The formulas hold for every polarization.
MARSIS is a nadir looking sounder which employs a linear polarization. Taking $\theta_{i1} = 0$, the expression for the reflectivity coefficient reduces to:

$$\rho = \frac{\sqrt{\epsilon_1} - \sqrt{\epsilon_2}}{\sqrt{\epsilon_1} + \sqrt{\epsilon_2}} \tag{1.46}$$

Hence, in this case it is not possible to distinguish between vertical and horizontal polarization for a nadir-looking geometry.

1.3.2 Dielectric properties of materials

As shown before, the measure of the dielectric permittivity of a target allows to estimate the morphology and the properties of the target itself. However, the EM behavior of natural materials is non-trivial. Many dielectric losses, related to different mechanisms, can occur at several frequencies, which makes the dielectric permittivity frequency dependent. Dielectric relaxation represents one of the primary mechanisms by which electromagnetic energy is absorbed in natural materials, resulting in frequency-dependent permittivity and loss. This phenomenon occurs when the polarization mechanisms within a material cannot instantaneously follow the oscillating electric field of the radar wave. The relaxation effects are present at specific frequency ranges, due to the different polarization phenomenon (i.e., electronic, ionic, dipolar) that originates them. Here we report the main dielectric properties of materials which are of interest for planetary radar observations, specifically MARSIS and ground planetary radar antennas.
As it will be described in detail in Chapter **3** MARSIS subsurface penetrating capabilities have been tailored for detecting the presence of liquid water in Martian terrain. Liquid water has a very high dielectric constant ($\epsilon' \approx 80$ at MHz frequencies) compared to common planetary minerals, making water-rich interfaces highly reflective to radar waves [9]. This high permittivity creates a strong dielectric contrast when liquid water is present beneath ice or regolith, resulting in bright radar reflections that can be readily detected even at depth.
The EM properties of water ice are considerably different. As shown in Figure **1.4**, pure water ice has a real dielectric constant of approximately $\epsilon' \approx 3.2$, which remains essentially constant from 10 MHz to over 300 GHz and shows only weak temperature dependence. The loss tangent of pure ice at cold temperatures typical of Mars (< 250 K) is very low ($\tan\delta < 0.001$), making ice nearly transparent to radar signals and allowing for deep penetration [10]. However, the presence of impurities or salts in ice can significantly increase the loss tangent and reduce penetration depth. Laboratory measurements have shown that even small concentrations of dissolved salts can increase radar attenuation by orders of magnitude, particularly at lower frequencies [11].
The dielectric properties of Martian regolith vary significantly across the planet's surface and with depth, reflecting variations in composition, porosity, and volatile content. Figure **1.5** shows the Mars global permittivity map obtained from MARSIS orbital radar measurements [12]. These observations reveal systematic variations with latitude: tropical regions typically show higher permittivity values ($\epsilon' \approx 6 - 10$), while higher latitudes exhibit lower values ($\epsilon' \approx 2 - 5$). This

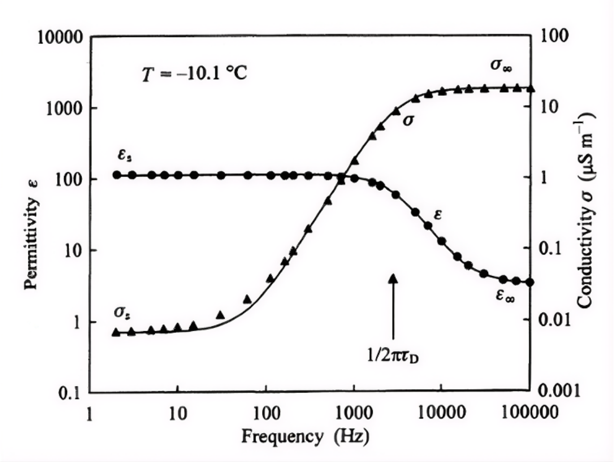

Figure 1.4: Real part of permittivity and conductivity as a function of frequency (2 Hz-100 kHz) of a single crystal of pure water ice measured at 263 K. From [5].

latitudinal gradient has been interpreted as evidence for increasing water ice content in the shallow subsurface at higher latitudes.

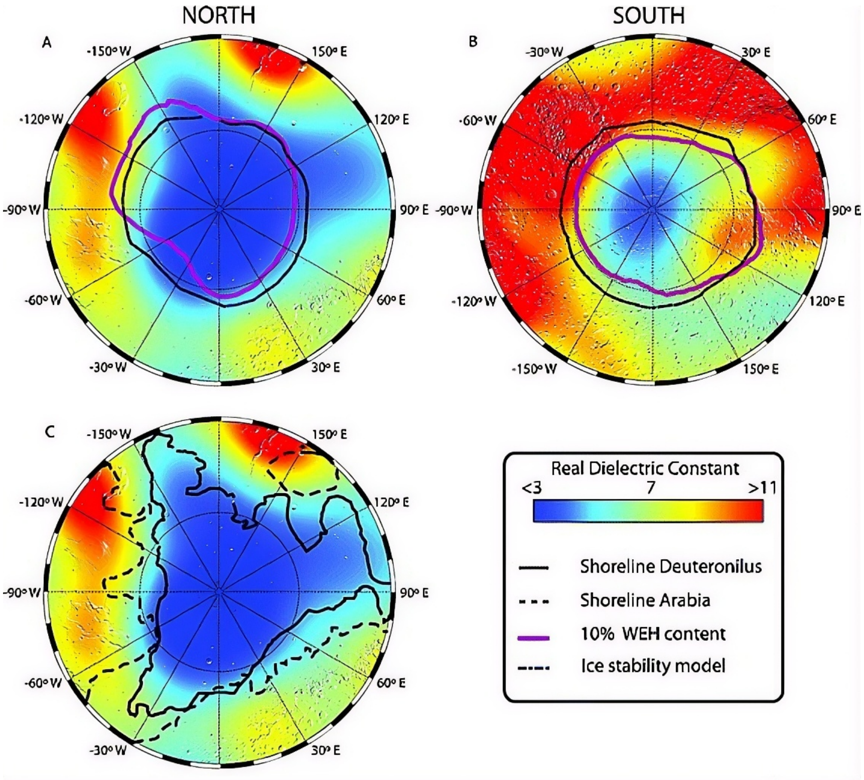

Figure 1.5: Maps of the Martian (**a**, **c**) northern and (**b**) southern hemisphere, from 30°S to the pole, displayed in polar stereographic projection. The blue-red colors represent the surface dielectric values measured by MARSIS and interpolated by spherical harmonics of order 8. Base map is a Mars Orbiter Laser Altimeter (MOLA) shaded-relief image. Low values (blue) are best explained by low density materials and/or presence of ice, while high values (red) indicate the presence of higher density volcanic materials. From [12].

The permittivity map generated from complementary observations from SHARAD at higher frequencies (20 MHz) shows a behavior that reflects Mars hemispherical dichotomy, with highland regions exhibiting $\epsilon' \approx 7-10$ and lowland regions showing $\epsilon' \approx 3-7$ [13]. These differences likely reflect variations in both composition (more weathered, porous material in the highlands) and the degree of cementation or induration.

These orbital measurements have been confirmed by subsequent in-situ observations from rover-mounted ground-penetrating radar. Data from the Zhurong rover's RoPeR instrument at the Tianwen-1 landing site in Utopia Planitia indicate permittivity values of $\epsilon' \approx 3-4$ in the upper 5 meters, with loss tangent values ranging from 0.006 to 0.011 [14]. These values are consistent with a dry, porous regolith composition of the Mars surface and uppermost subsurface.

Several laboratories studies have been conducted on Martian analog materials to investigate the frequency and temperature dependence of EM of these samples [9]. Their measurements, spanning 1 MHz to 1 GHz and temperatures from $180-300$ K, reveal that magnetic minerals such as magnetite and hematite, that are likely present on Mars, can exhibit significant dielectric relaxation losses at MARSIS frequencies, particularly at warmer temperatures. Most importantly, the amount of water absorbed by the regolith greatly influences the measured value of the dielectric permittivity.

While the previous discussion focused on Martian materials relevant to MARSIS observations, the dielectric properties of asteroid regolith are equally relevant to the ground-based radar observations presented in Chapter **4**. For ground-based planetary radar observations of asteroids at GHz frequencies (considering Arecibo and Goldstone), the dielectric properties of regolith materials are primarily governed by composition and porosity. Most dry, nonmagnetic, silicate minerals show relatively low permittivity values ($\epsilon' \approx 3-7$) with minimal frequency dependence across the MHz-GHz range [15]. Hence, the value of the dielectric permittivity of regolith at GHz frequencies can be considered constant.

1.4 Surface scattering

Until now we have described the interaction between an EM plane wave and a perfectly smooth surface according to the Fresnel reflection laws. The situation becomes more complicated when we consider a real radar planetary observation: the radiation pattern of a transmitting parabolic antenna generates a beam of waves with a spherical wavefront that impacts a generally rough planetary surface. In this scenario, the incident wave is partially reflected in the specular direction (with $\theta_i = \theta_r$) and partially scattered in all directions. A monostatic radar (transmitter and receiver are at the same location) receives the backscattered component of the scattered energy. Thus, theoretically, a monostatic radar would receive no return power from a completely smooth (specular) surface except for normal incidence (see Figure **1.6**). In this scenario, the angular radiation pattern of the reflected wave is a delta function centered about the specular direction. The specular component is often referred to as the coherent scattering component, while the scattered component is known as the diffuse or noncoherent component. As the surface becomes rougher, the coherent component becomes more negligible, until the radiation pattern approaches the theoretical limit of a *Lambertian* surface, defined as a surface whose radiance is independent of the angle of observation [16] or, equivalently, a surface that radiates isotropically.

1.4.1 Surface parameters

The degree of roughness of a surface depends on the wavelength of the incident wave. Thus, a given surface may appear very rough at optical wavelength but very smooth at microwave frequencies.

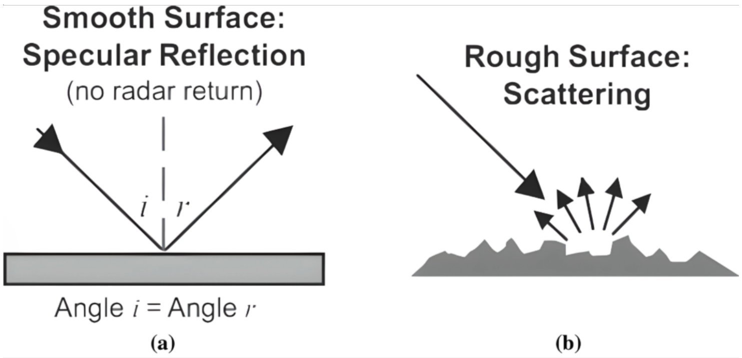

Figure 1.6: **b**: sketch of a specular reflection from a smooth surface. **b**: scattering from a rough surface. From [17].

Before digging into the statistical theory of random rough surfaces we introduce some hypothesis that are considered valid a priori for the following analysis. A random surface is *isotropic* if the statistics of that surface are independent of direction along the surface. This assumption often simplifies the mathematics of wave scattering, since any vector representation can be replaced by its modulus. There are more sophisticated mathematical descriptions that takes into account for the anisotropy of the surface of interest, since natural surfaces can be anisotropic if the processes forming them are themselves directionally dependent. Take as an example the sea surface, where wind direction can generate a strongly directional surface.

In this section, we'll assume that the probability of a surface point being at a particular height $p(z)$ is independent of z. More in general, any statistical properties dependent on two or more surface points depend only on the vector joining them and not on their absolute positions. The surface is said to be *stationary*, such that the statistics of one section of the surface will be the same as the statistics determined from a different section of the same surface. This assumption, although extremely convenient, holds only if the lengths of the sections are all the same. As a matter of fact, if the length is altered then the statistical properties will change accordingly.

The surfaces are said to be *ergodic* if any statistical average taken over many different parts of one surface realization (spatial averaging) is the same as an average over many realizations (ensemble averaging) [18]. The ergodicity of a surface holds under the same conditions of stationary surfaces. Hence, surfaces that are ergodic must be stationary. The concept can be extended also to rough surfaces with time varying profiles, such as, again, sea surfaces: averaging over the same part of the surface at many different times (temporal averaging) will lead to the same statistical description as spatial or ensemble averaging.

A rough surface is usually described in terms of its deviation from a smooth "reference surface". In particular, the two fundamental parameters commonly used to characterize surface roughness are the *height probability density function p(z)* and the *surface correlation function* $\rho(\epsilon)$, which accounts for the spread of heights about the reference surface and the variation of these heights along the surface.

Figures **1.7a** and **1.7b** show a random isotropic surface and its corresponding height profile, respectively. The deviation of a surface from the smooth reference surface is represented by the function $z(\mathbf{x})$, where z is the height of the surface from the reference surface and $\mathbf{x}$ is the position vector of points on the reference surface. If we take the mean surface $<z>=0$ as the reference surface, we can rearrange the height profiles into height bins (Figure **c**). The distribution of surface heights is described by the height probability density function $p(z)$, which is Gaussian for a random surface [19]:

$$p(z) = \frac{1}{\sqrt{2\pi\sigma^2}} e^{-z^2/2\sigma^2} \tag{1.47}$$

where σ is the height standard deviation:

$$\sigma =< z^2 >^{1/2}= \left[\int_{-\infty}^{\infty} z^2 p(z) dz\right]^{1/2} \tag{1.48}$$

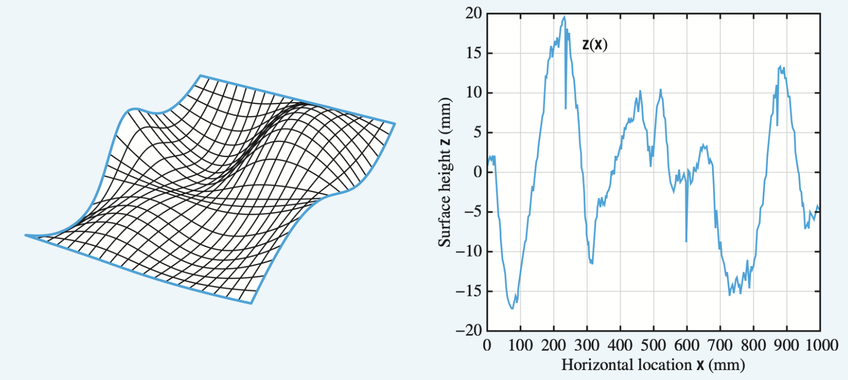

(a) Random surface

(b) Height profile

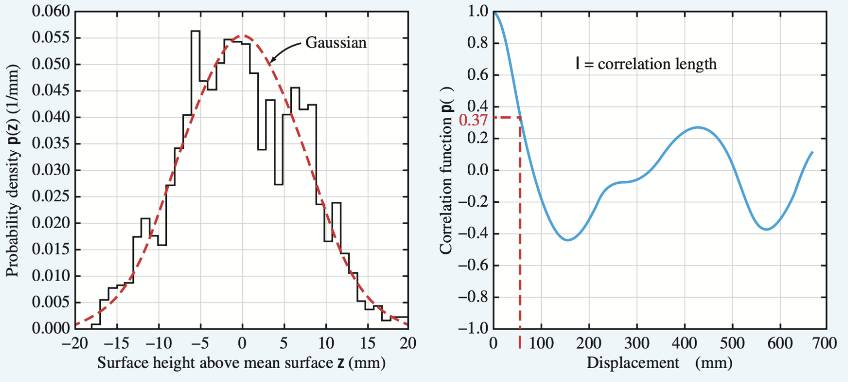

(c) Histogram

(d) Correlation function

Figure 1.7: Random, isotropic surface $z(x, y)$: **a**: pictorial view, **b**: measured height profile $z(x)$, **c**: pdf of digitized height profile, and **d**: autocorrelation function $\rho(\epsilon)$, where ϵ is the displacement between two points on the surface. From [4].

The surface correlation function $\rho(\epsilon)$ is a measure of correlation between the surface at two different locations, say (x, y) and (x', y'):

$$\rho(\epsilon) = \frac{< z(x, y) z(x', y') >}{\sigma^2} \tag{1.49}$$

with ϵ being the lateral separation between the two locations. As shown in panel **d**, the correlation function decreases with ϵ until a certain value, called *correlation length* l, such that, for distances greater than l, the heights at the two locations are considered statistically uncorrelated. Quantitatively, the correlation length is defined as:

$$\rho(\epsilon) = e^{-1} \quad \epsilon = l \tag{1.50}$$

The theory of wave scattering from rough surfaces often assumes two forms for the correlation function, namely the exponential $\rho_c(\epsilon)$ and the gaussian correlation function $\rho_g(\epsilon)$:

$$\rho_c(\epsilon) = e^{|\epsilon|/l} \tag{1.51}$$

$$\rho_g(\epsilon) = e^{-\epsilon^2/l^2} \tag{1.52}$$

We remark that, in case of a perfectly smooth surface, $l = \infty$.
Clearly the statical definition of a random surface has some intrinsic limitations. For instance, some surfaces cannot possess Gaussian height distributions because of their inherent asymmetry (i.e. abrasive processes like rain erosion skew surfaces toward negative heights). Also, surface correlation functions should not be regarded as reliable for distances greater than about 1/10 of the surface extent [20]. In addition to that, surface correlations sometimes appear to contain more than one correlation length [21], especially when more than one process is responsible for forming the surface. In practice rough surfaces are unlikely to possess correlation functions that are as smooth as those given by Equations **1.51** and **1.52**. That's why there are statistical techniques more appropriate for describing random surfaces as fractals.
When we deal with planetary radar observations, the problem with rough surfaces measurements arises from the inherent nature of these surfaces: roughness exists over a wide range of scales, hence there's no common reference for the target surface horizontal scale. Therefore, the method used to measure the characteristics of a random rough surface must be carefully chosen according to the application for which the surface is of interest.

1.4.2 The Rayleigh criterion

Usually, the term "smooth surface" is defined as a specular surface that obeys the Fresnel reflection and transmission laws. A more empirical and widely used definition is that a surface may be considered smooth if the observed scattering coefficient or emissivity is in close agreement with values calculated when assuming the surface to be specular [15]. From a measurement standpoint, in case of a nadir-looking planetary radar, a surface may be considered specular if the backscattered power for off-nadir angles is almost negligible. If the surface is made slightly rough, the magnitude of the coherent nadir component decreases slightly as well, and the angular pattern of the backscattered coefficient $\sigma(\theta)$ will assume an exponentially decaying shape with θ for the angular region close to nadir.

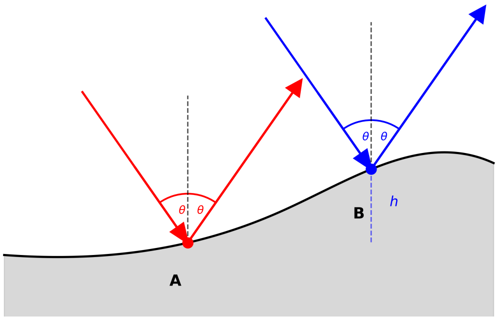

Figure 1.8: Diagram for determining the phase difference between two parallel rays scattered from different points.

Consider a monochromatic plane wave incident at some angle θ on a surface like in Figure **1.8**.

The *Rayleigh criterion* states that if the phase difference $\Delta\theta$ between the two reflected rays is less than $\pi/2$ radians, then the surface may be considered smooth. Mathematically, the phase difference can be written as:

$$\Delta\theta = 2kh\cos\theta \tag{1.53}$$

if we set $\Delta\theta < \pi/2$:

$$h < \frac{\lambda}{8\cos\theta} \tag{1.54}$$

The same formula can be applied to a random surface characterized by a standard deviation height σ by replacing it with h. While the Rayleigh criterion provides a useful first-order classification of surface smoothness, it proves insufficient for modeling scattering from natural surfaces at microwave frequencies, where the wavelength is typically comparable to the surface RMS height σ. A more stringent criterion is therefore required. To address this limitation, we adopt the criterion originally developed to define the far-field distance of an antenna aperture. This criterion stipulates that the maximum phase difference between rays originating from the center and the edge of the aperture must not exceed $\pi/8$ radians. Applying this condition, known as the *Fraunhofer criterion*, to surface scattering yields:

$$\sigma < \frac{\lambda}{32\cos\theta} \tag{1.55}$$

The above criterion appears to be more consistent with experimental observations [15].

1.4.3 Surface scattering models

Several models for EM waves scattering from random rough surfaces have been developed during the last two centuries. Such models can be categorized into empirical models, analytical models and a combination of the two. Although empirical models are the simplest, they depend directly on the experimental conditions, while theoretical models, even if more complex, can give a significant understanding of the physics of the interaction [22]. Theoretical models could not lead to an exact solution of the equation governing the process analyzed, hence approximate methods have to be used. In this section we provide an overview of two of the main analytical models developed in the last century: the *Kirchhoff Approximation* (KA) and the *Small Perturbation Method* (SPM).

The Kirchhoff Approximation, also known as the Physical Optics (PO) approximation or tangent plane approximation, treats the scattering problem by assuming that each point on a rough surface scatters incident radiation as if it were part of an infinite plane tangent to the surface at that point [4]. This approach reduces the complex problem of scattering from an irregular surface to a superposition of reflections from locally flat facets, with the scattered field computed by integrating contributions from all illuminated surface elements. The validity of the KA is based on two primary conditions. First, the radius of curvature of the surface must be large than the wavelength, ensuring the validity of the tangent plane approximation. Second, the surface slopes must be relatively gentle, typically requiring that the RMS slope be less than approximately 0.3 radians [15]. These conditions are mostly satisfied for surfaces that are smooth at the scale of the wavelength but may exhibit significant height variations, making KA particularly well-suited for modeling scattering from surfaces with large-scale undulations and moderate roughness. In the context of planetary radar observations, this regime corresponds to surfaces where $k\sigma \ll 1$ and $kL \gg 1$, where k is the wavenumber, σ is the RMS height, and L is the correlation length of the surface roughness.

The principal advantage of the Kirchhoff Approximation lies in its relative computational simplicity and its ability to handle surfaces with arbitrary height variations, provided the slope constraints are satisfied.

It naturally accounts for shadowing effects, which occur when parts of the surface are hidden from either the transmitter or receiver by intervening topography. This becomes increasingly important at low grazing angles.

The KA also provides a good description of the essential physics of quasi-specular scattering, where the backscattered return is dominated by surface facets oriented perpendicular to the radar line-of-sight (LOS). This makes it particularly valuable for modeling orbital radar observations of planetary surfaces, where near-nadir geometries often satisfy the gentle-slope requirement [23]. However, KA fails when surface slopes become too steep or when the surface exhibits fine-scale roughness comparable to the wavelength. In these cases, multiple scattering between neighboring surface elements becomes significant, and the independent facet assumption breaks down. The approximation also does not account for diffraction effects that arise when the surface contains sharp edges or abrupt discontinuities.

The Small Perturbation Method approaches the scattering problem from the opposite side, treating the rough surface as a small perturbation to a perfectly flat interface. The scattered field is expressed as a perturbation series, with successive terms accounting for higher-order scattering effects. In practice, only the first-order term is typically retained, yielding what is known as the first-order SPM or Bragg scattering model [24, 25]. The validity of first-order SPM requires that surface height variations be small compared to the wavelength ($k\sigma < 0.3$) and that surface slopes be gentle ($kL > 6$, approximately) [4]. Under these conditions, the scattered field results primarily from constructive interference of waves scattered from surface features. This means that the backscattered signal is sensitive to surface roughness components with wavelengths of approximately half the radar wavelength in the direction of observation.

The primary strength of the SPM lies in its ability to predict the angular and polarimetric characteristics of scattering from slightly rough surfaces with good accuracy. Unlike the Kirchhoff Approximation, SPM explicitly accounts for the coupling between incident and scattered fields through the surface boundary conditions, making it particularly well-suited for analyzing polarimetric radar observations. The method also naturally incorporates the frequency dependence of scattering from rough surfaces, predicting that backscatter increases with frequency for a given surface roughness, a behavior observed in many planetary radar datasets. The principal limitation of SPM is its restricted range of validity. For surfaces where $k\sigma$ approaches or exceeds unity, the perturbation series converges poorly or not at all, and the first-order approximation severely underestimates the scattered power. The method also fails to account for shadowing effects and cannot handle surfaces with significant height variations. In the context of planetary surfaces, which often exhibit roughness spanning multiple length scales, SPM is typically applicable only to the smoothest terrains or to specific roughness components when analyzing scattering from composite surfaces.

The Kirchhoff Approximation and Small Perturbation Method occupy complementary domains in the parameter space of surface roughness. KA excels for surfaces with large-scale undulations and moderate slopes, while SPM applies to surfaces with fine-scale roughness but small height variations. Real planetary surfaces often exhibit roughness at multiple scales, falling outside the strict validity range of either model. This has motivated the development of hybrid approaches, such as the Integral Equation Method (IEM) [26], which combines elements of both KA and SPM to extend the range of applicability. Nevertheless, understanding the physical principles and limitations of these foundational models remains essential for interpreting radar observations and for assessing when more sophisticated numerical methods may be required.

1.4.4 Scattering regimes

When a radar system transmits an electromagnetic wave towards a planetary surface or small body, different points within the illuminated area contribute to the total reflection with discrete sinusoidal components characterized by different phases ϕ_i, which depend on each point's range from both transmitter and receiver. The total path length of the i-th contribution corresponds to the distance from the transmitter to the generic scattering point P_i plus the distance from that point to the receiver. In monostatic configurations, such as MARSIS observations, this total range is called the two-way range, representing the same distance traversed twice in opposite directions by the electromagnetic wave. If a sinusoidal wave with zero initial phase is transmitted through vacuum, scattered by point P_i, and eventually received, its final phase can be expressed as:

$$\phi_i = 2\pi f \Delta t_i = \frac{2\pi}{\lambda} l_i = k l_i \tag{1.56}$$

where Δt_i is the time interval between transmission and reception, l_i is the total path length, and $k = 2\pi/\lambda$ is the free-space wavenumber.

For nadir-looking subsurface sounding radars such as MARSIS, where the long wavelength (54–167 m, corresponding to frequencies of 5.5–1.8 MHz) enables penetration into the Martian subsurface, a reflector may be located beneath the surface. In this case, the signal propagates partially through material with dielectric constant ϵ, and the total phase becomes:

$$\phi_i = 2k(r^{(i)}|_s + \sqrt{\epsilon} r^{(i)}|_{ss}) \tag{1.57}$$

where $r^{(i)}|_s$ and $r^{(i)}|_{ss}$ represent the one-way path segments above and below the surface, respectively.

In bistatic planetary radar observations, such as those employing the Goldstone Solar System Radar (DSS-14, 70-m dish) or DSS-63 antenna (70-m) at the Madrid Deep Space Communications Complex as transmitters, with European radio telescopes (Medicina, Effelsberg, Noto) serving as receivers, the geometry differs significantly. The received phase is:

$$\phi_i = k(r_1^{(i)} + r_2^{(i)}) \tag{1.58}$$

where $r_1^{(i)}$ and $r_2^{(i)}$ are the ranges from transmitter to scatterer and from scatterer to receiver.

In both geometries, sinusoidal contributions from different locations within the effective reflecting area add coherently. The first *Fresnel zone* is defined as the surface area from which scattered waves differ in phase by less than π, corresponding to a path length difference of $\lambda/2$. Waves scattered from within this zone interfere constructively, determining the primary contributing area to the specular return. The Fresnel zone is circular for monostatic near-nadir geometries and elliptical for bistatic or off-nadir configurations (Figure **1.9**).

For monostatic subsurface radar observations, assuming a spherical wavefront and a planar subsurface interface, the radius of the first Fresnel zone within a medium of refractive index $n = \sqrt{\epsilon}$ is:

$$r_{FZ} = \sqrt{\frac{\lambda}{2}\left(h + \frac{d}{n}\right)} = \sqrt{\frac{\lambda h'}{2}} \tag{1.59}$$

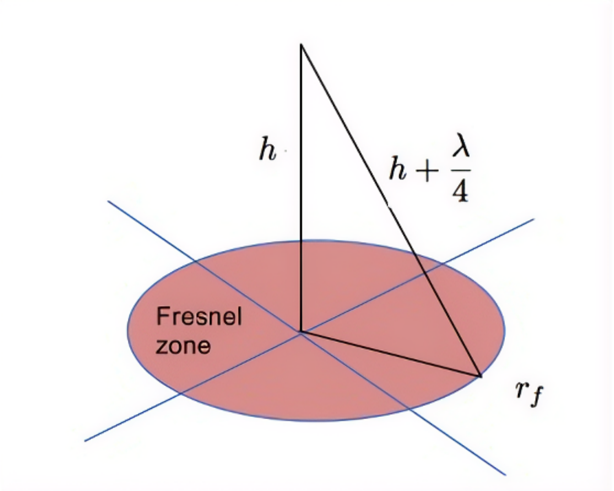

Figure 1.9: Geometry of the Fresnel zone

where $h' = h + d/n$ represents an effective altitude, that is the sum of the spacecraft altitude h above the surface and the subsurface target depth d. For MARSIS at typical Mars Express altitudes (250–800 km) and operating frequencies ($1.8 - 5.5$ MHz), the first Fresnel zone radius for surface reflections ranges from approximately 3 to 12 km. The assumption of interface planarity is justified since the orbital altitude remains more than an order of magnitude smaller than Mars' radius (~ 3390 km).

For bistatic observations of asteroids or planetary surfaces from Earth-based facilities, particularly for smaller bodies where target sphericity cannot be neglected, the first Fresnel zone forms an ellipse on the surface. The exact analytical expressions for the semi-major and semi-minor axes depend on the bistatic geometry (transmitter-target-receiver configuration), target radius, and incidence angle. For kilometer-sized asteroids observed from Earth at distances of $\sim 10^7$ km, the Fresnel zone can encompass a substantial fraction of the illuminated hemisphere.

As discussed in Section **1.4.1**, surface roughness causes the effective reflecting area to extend beyond the first Fresnel zone. However, this expansion is bounded by physical constraints specific to each radar configuration. For orbital sounders like MARSIS, the limit is determined either by antenna beamwidth or pulse resolution. For bistatic ground-based radar, the constraint arises from the geometric overlap between transmitter and receiver antenna footprints.

In bistatic radar configurations using transmitters such as Goldstone (operating at 8.56 GHz)) or DSS-63 (7.19 GHz with similar beamwidth) and European receiving stations, the observation is beam-limited when the overlap between transmitting and receiving antenna footprints is smaller than the potential scattering area determined by surface roughness. The footprint boundary is typically defined by projecting the half-power beamwidth (HPBW) onto the target surface. For a transmitter at distance r_1 with half-power beamwidth ϕ_{HPBW} (the full-width aperture angle where antenna gain drops by 3 dB relative to boresight, given by $\phi_{\text{HPBW}} \approx \lambda/D$ for a circular aperture of diameter D), the footprint area projects as an ellipse:

$$A_{\text{FP}} = \pi \left(\frac{r_1\, \phi_{\text{HPBW}}}{2} \right) \left(\frac{r_1\, \phi_{\text{HPBW}}}{2\cos\theta} \right) \tag{1.60}$$

where θ is the incidence angle. An analogous expression holds for the receiver footprint, replacing r_1 and ϕ_{HPBW} with r_2 and $\phi_{\text{HPBW,Rx}}$, respectively. For asteroid observations where both antennas are Earth-based and the bistatic angle is small, the effective footprint corresponds to the intersection

of two projected ellipses.
For pulsed radars like MARSIS, which possess range resolution through time-delay discrimination, rough surface scattering is ultimately limited by the pulse-limited footprint. As surface roughness increases and more energy scatters from off-nadir points, the instrument's range resolution effectively truncates the reflecting area to include only points whose range from the spacecraft lies within one resolution cell. Contributions from beyond this annulus appear as delayed returns and may be misinterpreted as subsurface features - a phenomenon known as *surface clutter*. Distinguishing genuine subsurface reflectors from off-nadir surface scatterers typically requires integration with topographic models derived from laser altimetry (e.g., MOLA for Mars [27]). The pulse-limited radius is obtained by replacing the Fresnel zone path difference ($\lambda/2$) with the range resolution $\Delta\rho = c/2B$, where B is the transmitted bandwidth:

$$r_{\mathrm{PL}} = \sqrt{2\Delta\rho h'} \tag{1.61}$$

For MARSIS, with 1 MHz bandwidth yielding 150 m free-space range resolution and effective altitudes of 250–800 km for surface observations, the pulse-limited footprint radius spans approximately 9–16 km. This significantly exceeds the Fresnel zone for most geometries, indicating that surface roughness often extends the effective scattering area well beyond the coherently-interfering Fresnel disk. The ratio between pulse-limited and Fresnel-limited footprints provides insight into the dominant scattering regime: when $r_{PL}/r_{FZ} \gg 1$, diffuse scattering dominates, and models accounting for roughness (such as those discussed in Section **1.4.3**) become essential for accurate interpretation of radar returns.

1.5 Radar Equation

Historically radar systems have been developed fo track identified targets, and to develop images of targets [28]. In both applications, radar performance depends on the strength of the signal received from the target and on the interference from unwanted signals.
Consider a generic radar system as the one showed in Figure **1.1**. The transmitter generates an EM wave with a certain power which is given in input to the antenna. If the antenna had an isotropic radiation pattern, the power density Q_i (watts per square meter) at distance R from the radiating antenna would be the total power divided by the surface area of a sphere of radius R:

$$Q_i = \frac{P_t}{4\pi R^2} \tag{1.62}$$

In practice the antennas have a directional beam pattern rather than an isotropic beam pattern. The transmitted power is concentrated into a region with a certain finite angular extent, with a width of several degrees in both the azimuthal and elevation directions (see Figure **1.10**). In this scenario, the power density at the center of the antenna beam pattern is higher than that from an isotropic antenna, since the transmitted power is focused on a smaller area. The *directivity* of an antenna is defined as the ratio between the power density for a lossless directional antenna and a hypothetical isotropic antenna. The *gain* G of an antenna is the *directivity* reduced by the losses the signal encounters while propagating through the medium between the antenna and the target. Hence, given G_t as the transmit antenna gain, we can rewrite Equation **1.62** as:

$$Q_i = \frac{P_t G_t}{4\pi R^2} \tag{1.63}$$

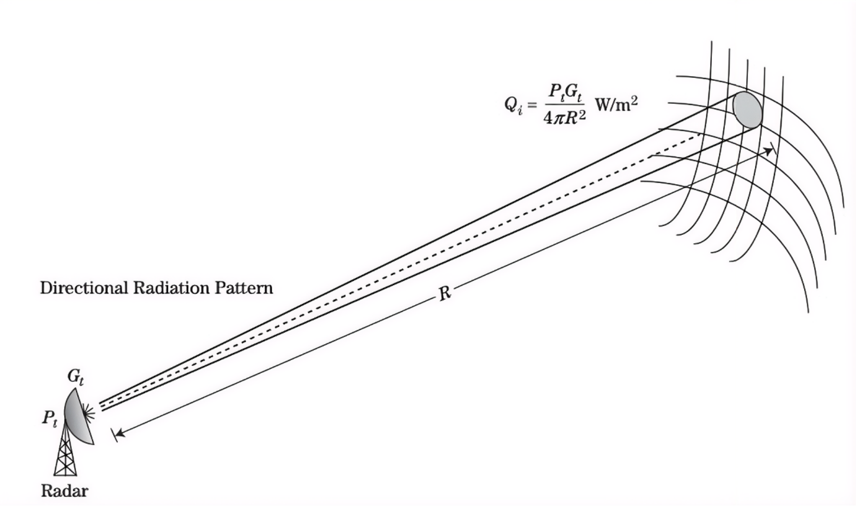

Figure 1.10: Power density at range R given transmit antenna gain G_t. From [2].

Once the transmitted signal hits the target, it induces time-varying currents, causing the target to re-radiate electromagnetic waves, part of which propagate back to the radar. The properties of the reflected wave collected by the receiver therefore depend on the physical size of the target, its shape, and the materials from which it is made. The *radar cross section* (RCS), represented with σ, is a measure of the power scattered in a given spatial direction when a target is illuminated by an incident wave. It can be interpreted as the effective area of an isotropic scatterer that would produce the same backscattered power density at the receiver as the actual target. Mathematically [29]:

$$\sigma = \lim_{R \to \infty} 4\pi R^2 \frac{|E^{\text{scat}}|^2}{|E^{\text{inc}}|^2} \tag{1.64}$$

where E^{scat} and E^{inc} are the scattered and incident electric field amplitudes respectively, R is the distance from the target, and σ has units of square meters.

In operational scenarios involving orbital radar observations of planetary surfaces, the radar cross section exhibits considerably greater complexity than idealized models suggest. The planetary surface is not uniformly illuminated, and the effective reflecting area, modulated by surface roughness, may be constrained either by the antenna beam pattern (beam-limited regime) or by the instrument's range resolution (pulse-limited regime). The reflected signal is not determined by geometry alone. Near-surface composition, where "near" refers to depths on the order of the radar penetration length, fundamentally scales the geometric radar cross section through the Fresnel reflectivity coefficient. Surface roughness operates across multiple length scales, each influencing the scattering mechanism differently. Large-scale roughness (comparable to or exceeding the Fresnel zone radius) modulates the spatial extent of the effective reflecting area, while small-scale roughness (comparable to the wavelength) redistributes energy from the coherent specular component into depolarized diffuse scattering. When the medium exhibits low electromagnetic losses, subsurface interfaces may generate additional reflections that interfere with the primary surface return. This layering of contributions from different depths and scattering mechanisms complicates echo interpretation, as individual components may be difficult to isolate without ancillary information or model-based inversion. Figure **1.11** illustrates the role of these mechanisms and their potential combinations in planetary radar returns. This complexity underscores both the interpretive challenges inherent to radar sounding analysis and the diagnostic potential of radio frequency remote sensing for characterizing surface morphology, composition, and subsurface structure across planetary bodies.

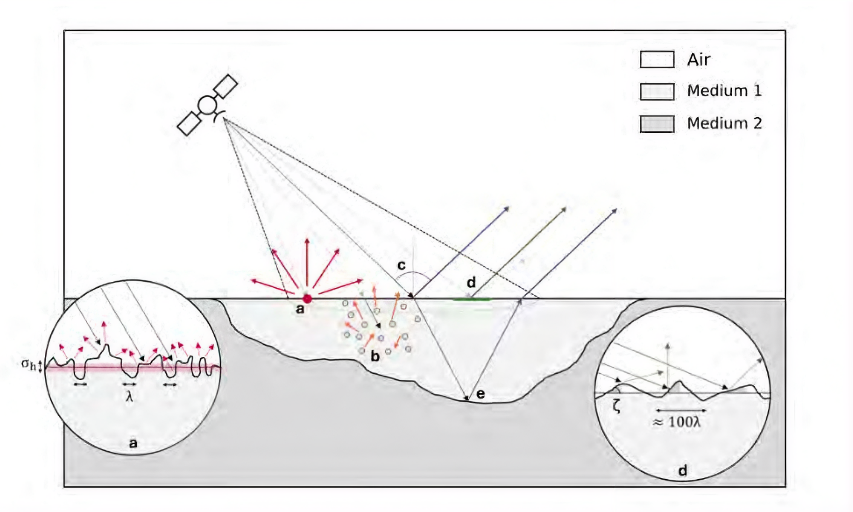

Figure 1.11: **a** diffuse scattering due to small-scale roughness. **b** Volume scattering due to subsurface contributions; **c** pure specular reflection from a smooth surface; **d** quasi-specular contribution due to large scale roughness and **e** specular contribution from potential dielectric interfaces in the subsurface. From [30].

The power reflected from a target toward the radar P_r can be expressed as:

$$P_r = Q_i \sigma = \frac{P_t G_t \sigma}{4\pi R^2} \tag{1.65}$$

The corresponding power density of the radio wave received back at the antenna is:

$$Q_r = \frac{P_r}{4\pi R^2} = \frac{P_t G_t \sigma}{(4\pi)^2 R^4} \tag{1.66}$$

The power P_r received by the antenna depends also on the antenna effective aperture A_e, which is related to the antenna gain according to:

$$G = \frac{4\pi \eta_a A}{\lambda^2} = \frac{4\pi A_e}{\lambda^2} \tag{1.67}$$

where η_a is the antenna efficiency. For a monostatic geometry, where the same antenna is used for both transmission and reception, the power received at distance R from the target can be written as:

$$P_r = \frac{P_t G^2 \lambda^2 \sigma}{(4\pi)^3 R^4} \tag{1.68}$$

while for a bistatic radar the formula becomes:

$$P_r = \frac{P_t G_t G_r \lambda^2 \sigma_{\text{bistatic}}}{(4\pi)^3 R_t^2 R_r^2} \tag{1.69}$$

where R_t and R_r denote the distance between the transmitter and the target and the target and the receiver, respectively. Note that σ_{bistatic} differs from σ, since the RCS depends also by the employed geometry and the radar wavelength.

In real applications, the signal received by the antenna is always corrupted by noise, a random interference with varying amplitude and phase. Noise can originate from both internal sources (e.g., thermal noise in the receiver electronics) and external sources (e.g., galactic background radiation). Thermal noise is particularly relevant as its power spectral density is approximately constant across all radar frequencies. The power P_n of the thermal noise in the radar receiver can be written as [31]:

$$P_n = kT_S B = kT_0 F B \tag{1.70}$$

where k is the Boltzmann's constant (1.38×10^{-23} watt-sec/K), T_0 is the standard temperature (290 K), T_s is the system noise temperature, B is the instantaneous receiver bandwidth in Hz, and F is the noise figure of the receiver subsystem (unitless). Usually, the signal-to-noise ratio (SNR) is what determines radar performance. It is defined as the received target signal power P_r divided by the noise power P_n. For a discrete target in a monostatic configuration:

$$\text{SNR} = \frac{P_t G^2 \lambda^2 \sigma}{(4\pi)^3 R^4 k T_0 F B} \tag{1.71}$$

There are several other sources that can reduce the power received by a radar antenna. The most common are due to clean air, rain, component losses, beam scanning, straddling, and several signal processing techniques. There are usually taken into account by introducing a term L_s (the so-called system loss) at the denominator of Equation **1.71**:

$$L_s = L_t L_a L_r L_{sp} \tag{1.72}$$

The transmit loss L_t considers the loss in the signal level as it travels from the transmitter to the antenna. The atmospheric loss L_a is due to the attenuation of the EM wave in the atmosphere as it travels from the radar to the target and back. Component losses when the signal travels from the antenna to the receiver, analogue to L_t, are represented by L_r. The term L_{sp} takes into account for all the effects due to signal processing, such as beam scan loss, straddle loss, automatic detection constant false alarm rate loss, and mismatch loss.

Chapter 2

Digital Signal Processing for Radar Sounding

Modern radar systems employ a variety of signal processing techniques to extract information from raw echoes, including signal integration to improve signal-to-noise ratio (SNR), pulse compression for fine range resolution, and angle-delay estimation for tracking. The development of increasingly more sophisticated algorithms has led digital signal processing (DSP) methods to displace analog processing in all contemporary radar systems. DSP encompasses a wide range of signal types, from acoustic waves to images. Radar signals, however, exhibit several distinctive characteristics compared to other signal classes, including their complex-valued representation, large instantaneous bandwidth, and the need to resolve both range and angle simultaneously. This Chapter describes the nature of radar signals and the processing techniques required to transform raw time-domain echoes into interpretable geophysical information. We begin with fundamental concepts in digital signal processing - sampling, Fourier analysis, and spectral estimation - and progressively build toward advanced imaging methods such as synthetic aperture processing, migration algorithms, and tomographic reconstruction.

2.1 Fundamentals of Signal Processing

2.1.1 Signals in Time and Frequency Domain

A signal can be represented either as a function of time, *s(t)*, or as a function of frequency, *S(f)*. These two representations are equivalent (related by the Fourier transform) and describe the same signal from complementary perspectives. The time-domain representation is more intuitive - a radar transmits a pulse at $t = 0$ and receives an echo at $t = 2R/c$, where R is the target range. On the contrary the frequency-domain representation reveals spectral content, bandwidth occupation, and filtering characteristics, which are critical for understanding range resolution, signal separation, and interference propagation [32].

The relationship between these domains is defined by the Fourier transform, a mathematical tool that decomposes a time-domain signal into its constituent sinusoidal components. Restricting consideration to continuous-time signals *s(t)* for the moment, the Fourier transform can be written as:

$$S(f) = \mathcal{F}\{s(t)\} = \int_{-\infty}^{\infty} s(t)e^{-j2\pi ft}dt \tag{2.1}$$

where f denotes frequency. The time-domain signal can be regained by taking the inverse Fourier transform as:

$$s(t) = \mathcal{F}^{-1}\{S(f)\} = \int_{-\infty}^{\infty} S(f)e^{j2\pi ft}df \tag{2.2}$$

the exponential kernel $e^{j2\pi ft} = \cos(2\pi ft) + j\sin(2\pi ft)$ is interpreted as a basis function, decomposing the signal into sinusoidal components of different frequencies. This duality between time and frequency is fundamental in signal processing: Table **2.1** summarizes the main properties of the Fourier transform. The convolution property is necessary to matched filtering (see Section **2.2.4**), where time-domain convolution of the received signal with a matched filter template becomes simple multiplication in the frequency domain. The time-shifting property explains how target range manifests as a linear phase term in the frequency domain, while the scaling property relates pulse duration to bandwidth. Parseval's theorem states that signal energy is conserved between domains, a critical consideration for power budget calculations [32].

Property	**Time Domain**	**Frequency Domain**
Linearity	$a\,s_1(t) + b\,s_2(t)$	$a\,S_1(f) + b\,S_2(f)$
Time Shifting	$s(t - t_0)$	$S(f)\,e^{-j2\pi ft_0}$
Frequency Shifting	$s(t)\,e^{j2\pi f_0 t}$	$S(f - f_0)$
Time Scaling	$s(at)$	$\frac{1}{\lvert a\rvert} S\left(\frac{f}{a}\right)$
Convolution	$s_1(t) * s_2(t)$	$S_1(f)\,S_2(f)$
Multiplication	$s_1(t)\,s_2(t)$	$S_1(f) * S_2(f)$
Parseval's Theorem	$\int_{-\infty}^{\infty} \lvert s(t)\rvert^2\,dt$	$\int_{-\infty}^{\infty} \lvert S(f)\rvert^2\,df$

Table 2.1: Main properties of the Fourier transform.

A fundamental limitation in signal processing that arises from the Fourier theory is that a signal cannot be arbitrarily localized in both time and frequency simultaneously. This condition, known as the *uncertainty principle* [33], can be written as:

$$\Delta t \cdot \Delta f \geq \frac{1}{4\pi} \tag{2.3}$$

where Δt and Δf represent the duration and the bandwidth of the signal, respectively. In practice, this inequality means that signals of short duration have wide spectral width and vice versa. As we'll see in the following sections of this Chapter, small Δt implies fine range resolution, while small Δf signifies good Doppler resolution.

Real-valued signals satisfy the Hermitian symmetry in the frequency domain, such that:

$$S(-f) = S^*(f) \tag{2.4}$$

where $*$ denotes complex conjugation. This symmetry implies that negative frequencies contain redundant information, making half of the available spectrum useless. To eliminate this redundancy and preserve phase information, essential for coherent radar processing, modern radar systems employ complex baseband representation [7].

Consider the analytic representation of a real passband signal, i.e. the analytic signal, defined as:

$$s_a(t) = s(t) + j\hat{s}(t) \tag{2.5}$$

where $\hat{s}(t)$ is the Hilbert transform of *s(t)*:

$$\hat{s}(t) = \frac{1}{\pi}\int_{-\infty}^{\infty}\frac{s(\tau)}{t-\tau}d\tau = s(t) * \frac{1}{\pi t} \tag{2.6}$$

The Hilbert transform is a linear operator that shifts all frequency components of a signal by $-\pi/2$ radians. Equivalently, in the frequency domain it corresponds to multiplication by $-j \cdot \text{sgn}(f)$, being $\text{sgn}(f)$ the sign function. The analytic signal built in this way has the property that its spectrum is zero for all negative frequencies: $S_a(f) = 0$ for $f < 0$, removing in this way the redundant component [34].

In practical applications, radar receivers do not compute the Hilbert transform directly. The complex signal representation is implemented by the *In-phase and Quadrature (I/Q) demodulation.* To illustrate this procedure, consider a real bandpass signal centered at carrier frequency f_c:

$$s(t) = A(t)\cos\left(2\pi f_c t + \phi(t)\right) \tag{2.7}$$

where *A(t)* is the time varying amplitude and $\phi(t)$ the phase of the signal. As shown in Figure **2.1**, the receiver splits the signal into two components. One channel, called the in-phase or "I" channel of the receiver mixes the received signal with an oscillator, called the local oscillator (LO), at the radar frequency. The other channel, called the quadrature phase or "Q" channel, mixes the signal with an oscillator having the same frequency but a 90° phase shift from the I channel oscillator. After low-pass filtering to remove high-frequency components terms, the resulting complex baseband signal is then:

$$I(t) = A(t)\cos\left(\phi(t)\right), \tag{2.8}$$

$$Q(t) = A(t)\sin\left(\phi(t)\right), \tag{2.9}$$

$$s_{\text{bb}}(t) = I(t) + jQ(t) = A(t)e^{j\phi(t)} \tag{2.10}$$

This down-conversion process removes the carrier frequency, yielding a complex-valued signal centered at DC (zero frequency), which can be sampled at a rate determined by the signal bandwidth B rather than the higher carrier frequency f_c [35].

2.1.2 Sampling

When a continuous analog radar signal passes through an analog-to-digital-converter (ADC) it is discretized into a finite number of points that represent the original waveform. Discretization in the independent variable, usually time, is called *sampling*. Specifically, a sampled signal s[n], is obtained by taking the value of the analog continuous signal $s_a(t)$ at multiples of some sampling interval T_s:

$$s[n] = s_a(nT_s), \quad n = -\infty, ..., +\infty \tag{2.11}$$

The Nyquist sampling theorem defines the minimum number of samples that must be taken to adequately represent the analog signal. Consider the signal $s_a(t)$ and its Fourier transform $S_a(f)$,

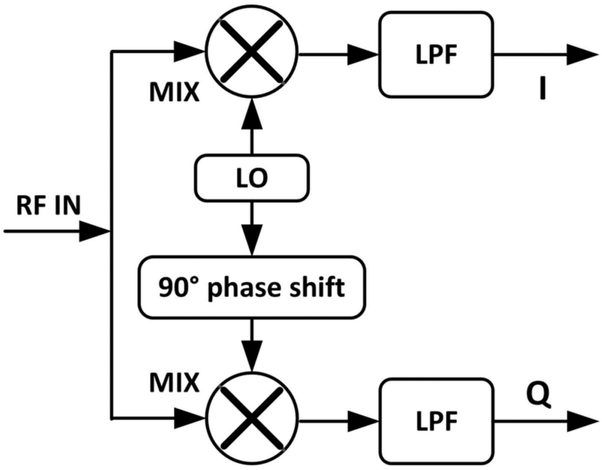

Figure 2.1: Theory of I/Q demodulation. MIX: RF mixer, LO: local oscillator, LPF: low-pass filters, I: in-phase components, Q: quadrature components. From [36].

assuming that it is band-limited to the interval $-B/2 \leq f \leq B/2$. The continuous-time signal $s_s(t)$ can be defined, using the shifting property reported in Table **2.1**, as:

$$s_s(t) = s_a(t)\left\{\sum_{n=-\infty}^{+\infty} \delta_D(t-nT_s)\right\} \tag{2.12}$$

$$= \sum_{n=-\infty}^{+\infty} s_a(nT_s)\delta_D(t-nT_s) \tag{2.13}$$

$$= \sum_{n=-\infty}^{+\infty} s[n]\delta_D(t-nT_s) \tag{2.14}$$

where $\delta_D(\cdot)$ is the Dirac delta function. From Equation **2.14**, it can be seen that $s_s(t)$ is a sequence of impulse functions occurring at the sample intervals nT_s, each weighted by a sample of the original analog signal.

Consider now the Fourier transform of $s_s(t)$. According to the theory of Fourier transforms, sampling in one domain causes a periodic replication of the signal in the complementary domain:

$$S_s(t) = \frac{1}{T_s}\sum_{n=-\infty}^{+\infty} S_a\left(f-\frac{k}{T_s}\right) = \frac{1}{T_s}\sum_{n=-\infty}^{+\infty} S_a(f-kf_s) \tag{2.15}$$

The spectrum $S_s(f)$ of the sampled signal consists of replications of the original spectrum every $1/T_s = f_s$ Hz. The original spectrum, and therefore the original signal, can be recovered as long as the replicas don't overlap. This is true if $s_a(t)$ is band-limited and if the sampling frequency satisfies the Nyquist criterion:

$$f_s = \frac{1}{T_s} > B \tag{2.16}$$

The minimum sampling rate of B samples/second is called Nyquist rate. When the sampling frequency is insufficient ($f_s \leq B$), the spectral replicas overlap, causing *aliasing*. High-frequency components of the original signal fold back into the baseband, appearing as false low-frequency components that cannot be distinguished from the original signal content. Mathematically, a frequency component at $f_0 > f_s/2$ will alias to an apparent frequency:

$$f_{\text{alias}} = |f_0 - k \cdot fs| \tag{2.17}$$

where k is chosen such that alias lies within the baseband interval [0, B/2]. Once aliasing occurs, the original signal cannot be recovered, as information has been permanently lost through spectral overlap [37].
The Nyquist sampling theorem is valid for each kind of signal, even if it is still at the carrier frequency or if I/Q demodulation is applied to form a complex-valued signal [28].

2.1.3 Quantization

The discretization of analog signals is called *quantization.* This is a necessary step in order to represent the amplitude of each sample with a finite number of bits in a digital hardware. A N-bit ADC maps the continuous range of input voltages to 2^N discrete output levels. If the ADC input ranges between $[-V_{\max}, V_{\max}]$, the quantization step size or the *least significant bit* (LSB) is:

$$\Delta = \frac{V_{\max}}{2^{N-1}} \tag{2.18}$$

Each sample is rounded to the nearest quantization level, which introduces a quantization error $\epsilon_q[n]$, such that $|\epsilon[n]| \leq \Delta/2$. It has been shown that for signals that span a significant portion of the ADC range, this error can be modeled as uniformly distributed white noise uncorrelated with the signal [38]. The quantization process therefore imposes a fundamental limit on the achievable signal-to-quantization-noise ratio (SQNR), which for a full-scale sinusoidal signal is approximately:

$$\text{SQNR} \approx 6.02N + 1.76 \quad \text{dB} \tag{2.19}$$

Each additional bit of resolution improves SQNR by approximately 6 dB [34].
The *dynamic range* of the ADC is defined as the ratio between the largest and smallest signals that can be represented. The theoretical dynamic range, expressed in decibels, is:

$$\text{DR} = 20 \log_{10} (2^N) \approx 6.02N \tag{2.20}$$

For radar applications, the choice of bit depth represents a trade-off between dynamic range and data volume. A larger N improves the ability to detect weak targets in the presence of strong clutter returns, but increases telemetry bandwidth and onboard storage requirements. MARSIS employs an 8-bit ADC for each of the I and Q channels, providing a theoretical dynamic range of approximately 48 dB [39]. This is enough for subsurface sounding applications where the echo dynamic range can be managed through automatic gain control (AGC) and onboard data compression.

2.1.4 Discrete Fourier Transform

The continuous Fourier transform introduced in Section **2.1.1** operates on infinite-duration signals defined over continuous time. In practice, radar systems process finite-length sequences of sampled data. The appropriate tool for spectral analysis of such sequences is the Discrete Fourier Transform (DFT), which converts a finite sequence of N complex samples in the time domain into N complex samples in the frequency domain. For a discrete-time signal $s[n]$ with $n = 0, 1, ..., N-1$, the DFT is defined as:

$$S[k] = \sum_{n=0}^{N-1} s[n] e^{-j2\pi kn/N}, \quad k = 0, 1, ..., N-1 \tag{2.21}$$

while the inverse DFT (IDFT) can be written as:

$$s[n] = \frac{1}{N} \sum_{n=0}^{N-1} S[k] e^{j2\pi kn/N}, \quad n = 0, 1, ..., N-1 \tag{2.22}$$

The DFT $S[k]$ is a finite, computable frequency domain representation of the signal $s[n]$. It represents the signal's frequency content at discrete frequency bins spaced by $\Delta f = f_s/N$, being f_s the sampling frequency. The k-th bin corresponds to frequency $f_k = k \cdot \Delta f = k \cdot f_s/N$. For complex-valued signals, the full frequency range $[-f_s/2, f_s/2]$ is represented, with bins $k = 0, ..., N/2 - 1$ corresponding to positive frequencies and bins $k = N/2, ..., N - 1$ corresponding to negative frequencies [34]. Several properties of the continuous Fourier transform, like linearity, time-shifting, and the convolution theorem are valid also for the DFT. However, it also introduces circular behavior: time-domain convolution becomes circular convolution, and shifts wrap around the N-point sequence.

While an N-point DFT is sufficient to represent an N-point discrete sequence, it is often preferable to increase the DFT resolution Δf. The frequency resolution of the DFT is given by:

$$\Delta f = \frac{f_s}{\mathrm{N}} \tag{2.23}$$

Δf can be increased in two ways: by increasing f_s or by increasing N. While the former requires hardware modification, the latter can be achieved through *zero-padding*. Zero-padding consists of appending zeroes at the end of the time-domain sequence until the desired length is reached. Since no new data is added, zero-padding does not modify the original signal - it only interpolates the DFT spectrum, providing a smoother representation of the underlying continuous spectrum. However, the fundamental frequency resolution remains limited by the observation time $T = M/f_s$, where M is the original signal length, which determines the mainlobe width of the spectral response [40].

In planetary radar applications, zero-padding improves visualization and peak-finding accuracy in Doppler spectra, and the frequency oversampling facilitates accurate frequency-domain filtering.

Direct computation of the DFT according to its definition requires $O(N^2)$ complex multiplications - N multiplications for each of the N output bins, which can be computationally prohibitive for typical radar datasets. The *Fast Fourier Transform* (FFT) [41] refers to any class of algorithms for efficient computation of the DFT, reducing the complexity to $O(N \log_2 N)$ operations. The first and most common used FFT is the radix-2 algorithm, which recursively decomposes the DFT of length $N = 2^m$ into two DFTs of length $N/2$, until reaching trivial 2-point transforms. Table **2.2** shows the computational cost of DFT compared to FFT for various number of samples. As typical of fast algorithms, the advantage over the DFT implementation grows with the number of samples.

Table 2.2: Computational cost comparison between DFT and FFT

N (samples)	**DFT operations**	**FFT operations**	**Speedup**
256	65,536	2,048	32×
1,024	1,048,576	10,240	102×
4,096	16,777,216	49,152	341×
16,384	268,435,456	229,376	1,170×

Modern radar processors universally employ FFT algorithms for pulse compression, Doppler processing, and SAR focusing. The requirement that N be a power of 2 (for radix-2 FFT) is typically met by zero-padding the data to the next power of 2.

The implicit assumption of the DFT is that the input sequence repeats periodically. When analyzing a segment of a longer signal, the abrupt truncation at the sequence boundaries introduces discontinuities, which manifest as spectral leakage, that is energy from a single frequency component spreads across multiple frequency bins, obscuring nearby components and degrading spectral resolution [42].

Consider a pure sinusoid at frequency f_0 sampled for duration T. If f_0 does not coincide exactly with a DFT bin frequency ($f_0 \neq k \cdot f_s/N$ for any integer k), the DFT exhibits a sinc-shaped spectrum centered at f_0, with sidelobes decaying as $1/f$. These sidelobes can mask weaker spectral components, a critical issue in radar when detecting small targets near strong clutter returns.

Windows are weighting functions applied to data to reduce the spectral leakage associated with finite observation intervals. Their function is to reduce the order of the discontinuity at the boundary of the periodic extension. A window function $w[n]$ is applied element-wise to the signal before the DFT:

$$s_w[n] = s[n] \cdot w[n] \tag{2.24}$$

Some of the most used window functions in radar applications are:

- **Rectangular window (no windowing)**
  - Definition:
  $$w[n] = 1 \tag{2.25}$$
  - Characteristics: narrowest mainlobe, highest sidelobes (approximately -13 dB).
- **Hann window (also called Hanning)**
  - Definition:
  $$w[n] = 0.5\left(1 - \cos\left(\frac{2\pi n}{N-1}\right)\right) \tag{2.26}$$
  - Characteristics: moderate mainlobe width, sidelobe level around -31 dB.
- **Hamming window** (Figure **2.2**)
  - Definition:
  $$w[n] = 0.54 - 0.46\cos\left(\frac{2\pi n}{N-1}\right) \tag{2.27}$$
  - Characteristics: similar to the Hann window, with slightly improved sidelobe suppression (approximately -43 dB).
- **Blackman window**
  - Definition:
  $$w[n] = 0.42 - 0.5\cos\left(\frac{2\pi n}{N-1}\right) + 0.08\cos\left(\frac{4\pi n}{N-1}\right) \tag{2.28}$$
  - Characteristics: wider mainlobe, excellent sidelobe suppression (approximately -58 dB).

The choice of window represents a trade-off between frequency resolution (mainlobe width) and dynamic range (sidelobe level). Narrower mainlobes improve the ability to distinguish closely-spaced frequency components, while lower sidelobes improve detection of weak signals near strong ones.

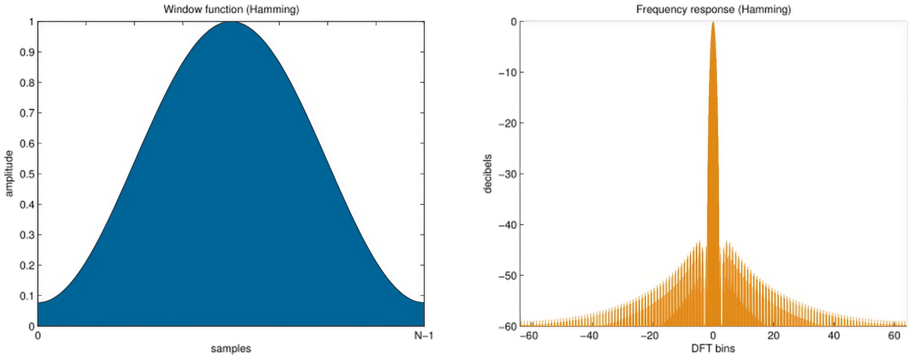

Figure 2.2: **left**: time-domain representation of the Hamming function. **right**: spectral representation.

2.1.5 Signal Integration

Integration is the process of combining together multiple samples of a signal, to obtain a single sample that has higher SNR than the individual samples. Two fundamentally different integration strategies exist, depending on whether phase information is preserved.

Coherent integration sums complex-valued samples while preserving phase relationships. Let's consider a measured signal s consists of a signal component, $Ae^{j\phi}$, and a noise component n. If we perform M independent observations of the signal, a sequence of measurements $s[n]$ can be formed. The noise is assumed to be independent and identically distributed. The coherently integrated signal is therefore:

$$s_{\text{coh}}[n] = \sum_{i=1}^{M} s_i[n] = \sum_{i=1}^{M}(Ae^{j\phi} + n_i[n]) = MAe^{j\phi} + \sum_{i=1}^{M} n_i[n] \tag{2.29}$$

The signal samples have been added all in phase with one another, hence the amplitude of the coherently integrated signal components becomes MA, and the signal power is $(NA)^2$. On the other hand, the noise power increases by a factor M, having assumed a white random noise. Therefore the SNR improves as:

$$\text{SNR}_{\text{coh}} = M \cdot \text{SNR}_{\text{single}} \tag{2.30}$$

This increase is called the integration gain. Coherent integration requires that the signal phase be stable across all observations, which imposes strict requirements on motion compensation and system phase stability [35].

When phase coherence cannot be maintained, due to target decorrelation, platform motion, or long integration times, incoherent integration can be reached either by summing the magnitude of the signal or its power (squared magnitude). The former is called noncoherent integration with a linear detector, while the latter uses a square law detector:

$$|s| = \sum_{i=1}^{M} |Ae^{j\phi} + n[n]| \tag{2.31}$$

$$|s|^2 = \sum_{i=1}^{M} |Ae^{j\phi} + n[n]|^2 \tag{2.32}$$

The signal power adds linearly (factor of M), but the noise power also adds linearly. If we consider Eq.**2.32**, the effective noise standard deviation increases only by a factor $\sqrt{M}$. The SNR improvement becomes:

$$\mathrm{SNR}_{\mathrm{inc}} = \sqrt{M} \cdot \mathrm{SNR}_{\mathrm{single}} \tag{2.33}$$

Noncoherent integration is less efficient than coherent integration. This result was expected since noncoherent integration doesn't take advantage of all of the information in the data. The trade-off is that incoherent integration is more robust and does not require phase stability [43].

MARSIS employs coherent integration for pulse compression and single-orbit SAR focusing, where phase stability is maintained, while incoherent integration is used when stacking echoes from multiple orbital passes where geometry variations destroy phase coherence. Similarly, ground-based bistatic observations of asteroids use coherent integration within individual observing sessions to construct delay-Doppler images, but require incoherent integration when combining data from multiple epochs separated by asteroid rotation or orbital motion [44].

2.2 Radar Waveforms and Range Processing

2.2.1 Continuous Wave (CW) Radar

The simplest radar configuration transmits a continuous, unmodulated sinusoidal waveform at a single carrier frequency f_c. Because of the continuous nature of CW emission, range measurement is not possible without some modifications to the radar operations and waveforms. However, CW radar remains the primary technique for ground-based planetary radar observations of small bodies, where Doppler spectroscopy - the analysis of frequency shifts in the received echo caused by the relative motion between the radar and the target - provides valuable information about target rotation and surface properties.

A CW radar emits a signal with constant frequency and constant amplitude, in the form:

$$s_{\mathrm{TX}} = A_{\mathrm{TX}} \cos\left(2\pi f_c t + \phi_0\right) \tag{2.34}$$

where A_{TX} is the transmitted amplitude and ϕ_0 is an arbitrary initial phase. When the incoming signal hits a target at distance R, the corresponding echo is received at a time delay $\tau = 2R/c$, where the factor of 2 accounts for the round-trip propagation in a monostatic configuration. Spectra of the radar echo from stationary targets and clutter will be concentrated at f_c. If there is a relative motion between the radar and the target, the center frequency will be shifted by an amount f_D, the Doppler frequency. The received signal can be expressed as:

$$s_{\mathrm{RX}} = A_{\mathrm{RX}} \cos\left(2\pi(f_c + f_D)(t-\tau) + \phi_0\right) \tag{2.35}$$

where A_{RX} accounts for propagation loss and target reflectivity.

By measuring the frequency difference between the transmitted and received signal it is possible to accurately extract target radial velocity. For a point target with radial velocity v_r (positive for receding, negative for approaching), the Doppler frequency is:

$$f_D = -\frac{2v_r}{\lambda} = -\frac{2v_r f_c}{c} \tag{2.36}$$

where the factor of 2 accounts for the two-way propagation path.

In case of extended targets such as asteroids, each surface element has different radial velocity due to rotation and irregular shape. The received signal is therefore a superposition of echoes from all illuminated surface facets, each contributing a Doppler-shifted component. Thus, the resulting Doppler spectrum contains information about the target's rotation period and surface roughness. Regions rotating toward the radar produce negative Doppler shifts (higher frequencies), while regions rotating away produce positive shifts (lower frequencies), creating a characteristic broadened spectrum whose width is proportional to the projected rotational velocity at the target's limbs.

Ground-based planetary radar facilities employ CW transmission for observing near-Earth asteroids and main-belt targets. The NASA Deep Space Network antennas DSS-14 (Goldstone, 70-m, X-band at 8.56 GHz) and DSS-63 (Madrid, 70-m, C-band at 7.19 GHz) serve as high-power transmitters, while radio telescopes such as the Green Bank Telescope (100-m) or Effelsberg (100-m) serve as sensitive receivers in bistatic configurations [45].

CW radars evaluate the phase difference $\Delta\phi$ between transmitted and received signals (Figure **2.3**), that is:

$$\Delta\phi = \frac{4\pi R}{\lambda} \tag{2.37}$$

At microwave frequencies phase measurements lead to high precision range determination (of the order of millimeters). However, it can be seen from Equation **2.37** that ranges differing by integer multiples of $\lambda/2$ are indistinguishable. This limitation is called *range ambiguity*: in case of a 10 GHz radar ($\lambda = 3cm$), the unambiguous range is only 1.5 cm. Thus, unmodulated CW radar can't be used for target range determination. Since the transmitted signal contains no time-varying features, there is no mechanism to determine the propagation delay τ. All targets illuminated simultaneously, regardless of range, produce echoes at the same frequency (plus Doppler), making them indistinguishable unless they have different velocities.

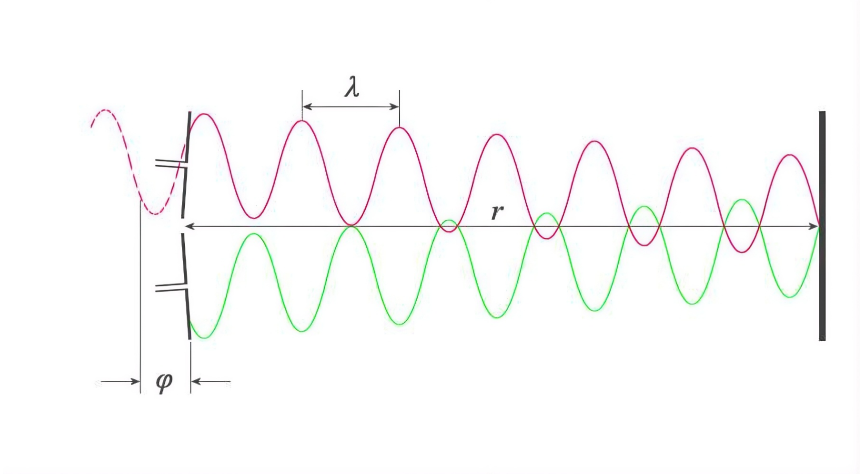

Figure 2.3: Phase difference between transmitted (red line) and received (green line) signals measured by CW radars. From [46].

Despite these limitations, CW radars are used in several applications, due to their relative cheap maintenance costs and high performances. Range measurements, if needed, can be achieved by changing the transmitting waveform, like the ones we present in the next sections. For planetary radar sounding, where both range resolution (to distinguish surface from subsurface layers) and velocity measurement (for orbital geometry) are essential, CW radar is insufficient. MARSIS transmits pulsed chirp waveforms to achieve fine range resolution (150 m) while maintaining coherent phase for Doppler processing.

2.2.2 Pulsed Radar

Pulsed radars emit short bursts or pulses and in the silent period collect the echo signals. In these instruments, unlike CW radars, the transmitter is turned off before the measurement is finished. This method offers a solution to the range ambiguity problem discussed before: by measuring the time delay between transmission and echo reception, pulsed radar can directly determine target range. This capability is crucial for applications like orbital subsurface sounding, where it is necessary to distinguish surface and subsurface reflections.

For a single pulse, the transmitted signal can be written as:

$$s_{\mathrm{TX}}(t) = A_{\mathrm{TX}} \cdot \mathrm{rect}\left(\frac{t}{\tau}\right) \cdot \cos\left(2\pi f_c t + \phi_0\right) \tag{2.38}$$

where $\mathrm{rect}(t/\tau)$ is the rectangular window function (unity for $|t| < \tau/2$, zero elsewhere) and τ is the pulse duration. The received echo is delayed by a factor $\tau_0 = 2R/c$ and it can be described as:

$$s_{\mathrm{RX}}(t) = A_{\mathrm{RX}} \cdot \mathrm{rect}\left(\frac{t-\tau_0}{\tau}\right) \cdot \cos\left(2\pi(f_c + f_D)(t-\tau_0) + \phi_0\right) \tag{2.39}$$

For a complete pulse train at pulse repetition interval T_{PRI}, multiple echoes are received:

$$s_{\mathrm{RX}}(t) = A_{\mathrm{RX}} \sum_{n=-\infty}^{\infty} \mathrm{rect}\left(\frac{t - nT_{\mathrm{PRI}} - \tau_0}{\tau}\right) \cdot \cos\left(2\pi(f_c + f_D)(t - nT_{\mathrm{PRI}} - \tau_0) + \phi_0\right) \tag{2.40}$$

Figure **2.4** shows an example of the pulse radar operation. T_{PRI} is defined as the time between successive pulses, while the rate at which pulses are transmitted is the pulse repetition frequency (PRF): $f_{\mathrm{PRF}} = 1/T_{\mathrm{PRI}}$.

The *duty cycle* D is the fraction of time spent transmitting:

$$D = \frac{\tau}{T_{\mathrm{PRI}}} = \tau \cdot f_{\mathrm{PRF}} \tag{2.41}$$

Typical radar systems have low duty cycles ($D \ll 1$), meaning that most of the time is dedicated to the reception of the signal. This approach is also energetically efficient, allowing high power transmission and low average power dissipation. Indeed, MARSIS pulses have a duration of $\tau \approx 250\mu$s and PRF ranging from 125 to 160 Hz according to the operational mode, yielding duty cycles of approximately $3-4\%$ [39].

The *range resolution* ΔR is the minimum separation in range between two targets that can be distinguished as separate. For a pulsed radar, range resolution is fundamentally limited by the signal bandwidth B, not the pulse duration:

$$\Delta R = \frac{c}{2B} \tag{2.42}$$

hence, fine range resolution requires large bandwidth. If we consider rectangular pulses like the ones in Figure **2.4**, the bandwidth is $B \approx 1/\tau$, implying that shorter pulses provide better range resolution. However, shorter pulses have low energy, meaning low SNR and limiting detection range. Pulse compression techniques (see Section **2.2.4**) have been developed to overcome this issue, allowing the use of long frequency-modulated pulses to achieve both high energy and fine range resolution.

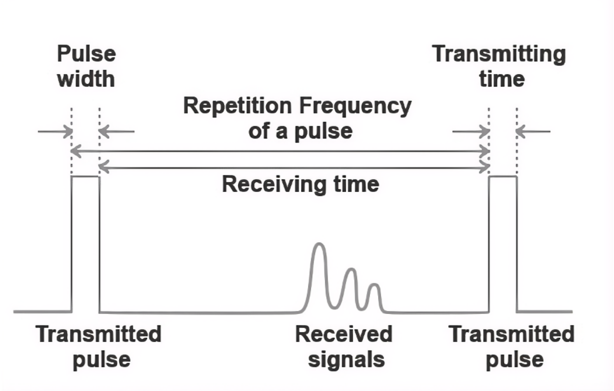

Figure 2.4: Example of a pulse radar signal. From [47].

While the range ambiguity discussed for CW radar is resolved in pulse radar, another issue arises: if a target echo is received after the transmission of the next pulse, it can be misinterpreted as a closer target. The range folding problem introduces an upper limit of target detection in pulse radar to avoid ambiguity. Specifically, the maximum *unambiguous range* R_{unamb} can be written as:

$$R_{\text{unamb}} = \frac{c \cdot T_{\text{PRI}}}{2} = \frac{c}{2 f_{\text{PRF}}} \tag{2.43}$$

For MARSIS, with $f_{\text{PRF}} = 127$ Hz for subsurface sounding, the unambiguous range is approximately ~ 1180 km, well above the maximum spacecraft altitude (~ 800 km).

If the Doppler frequency of the target is high enough to make an adjacent spectral line move inside the Doppler band of interest, the radar can be Doppler ambiguous. Therefore, in order to avoid Doppler ambiguities, radar systems require high PRF rates when detecting high speed targets. When a long-range radar is required to detect a high speed target, it may not be possible to be both range and Doppler unambiguous, since fine Doppler resolution requires high PRF, at the opposite of Equation **2.43**.

The *unambiguous Doppler frequency* is defined as half of the PRF:

$$f_{D,\text{unamb}} = \frac{f_{\text{PRF}}}{2} \tag{2.44}$$

In terms of target radial velocity:

$$v_{r,\text{unamb}} = \frac{\lambda \cdot f_{\text{PRF}}}{4} \tag{2.45}$$

For MARSIS at 4 Mhz ($\lambda = 75$ m), with PRF = 127 Hz, the unambiguous velocity is ~ 2.4 km/s, well above the orbital velocity variations encountered during typical observations.

2.2.3 Pulsed Chirp Radar

A linear frequency-modulated (LFM) pulse, or linear *chirp*, is characterized by a carrier whose instantaneous frequency increases (up-chirp) or decreases (down-chirp) linearly with time (see Figure **2.5** as an example). In case of an up-chirp, the transmitted signal can be written as:

$$s_{\text{TX}} = A \cdot \text{rect}\left(\frac{t}{\tau}\right) \cdot \cos\left(2\pi f_0 t + \pi \alpha t^2\right) \tag{2.46}$$

where α is the chirp rate (Hz/s), which indicates how rapidly the frequency sweeps. For an up-chirp starting at f_0 and sweeping to $f_0 + B$ over duration τ, the chirp rate is:

$$\alpha = \frac{B}{\tau} \tag{2.47}$$

this equation relates the pulse duration, the bandwidth and the chirp rate of a signal: a faster sweep (larger α) achieves a given bandwidth B in shorter time τ.

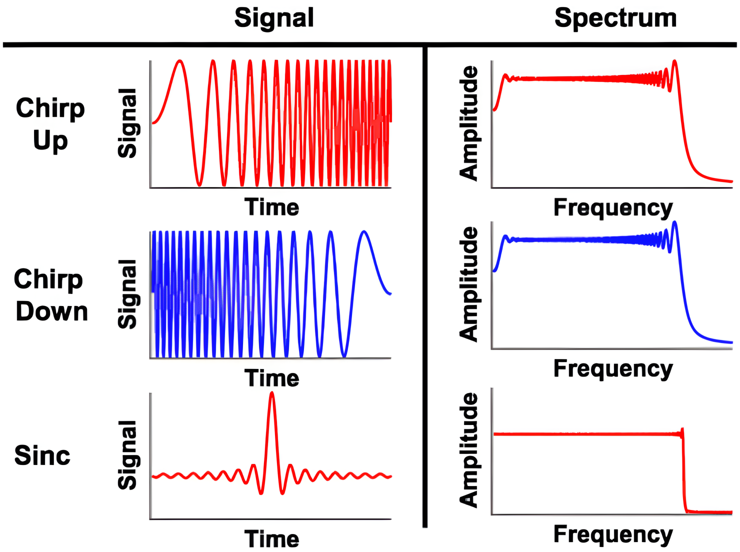

Figure 2.5: Chirp signal characteristics. From [48].

The *instantaneous frequency* f_{inst} describes the rate of oscillation at each moment in time. It is obtained by differentiating the total phase:

$$f_{\text{inst}}(t) = \frac{1}{2\pi}\frac{d\phi}{dt} = f_0 + \alpha t \tag{2.48}$$

which shows that the frequency increases linearly from f_0 at $t = 0$ to $f_0 + B$ at $t = \tau$.
The echo received by a target at distance R can be written as:

$$s_{\text{RX}} = A_{\text{RX}} \cdot \text{rect}\left(\frac{t - \tau_0}{\tau}\right) \cdot \cos\left(2\pi(f_0 + f_D)(t - \tau_0) + \pi\alpha(t - \tau_0)^2\right) \tag{2.49}$$

The quadratic phase term $\pi\alpha(t - \tau)^2$ carries the range information, which is extracted through matched filtering.
MARSIS transmits chirped pulses with the following nominal parameters [39]:

- pulse duration $\tau = 250\mu\text{s}$;
- bandwidth $B = 1$ MHz;
- PRF: typically 127 Hz;
- chirp rate: $\alpha = B/\tau = 4 \times 10^9$ Hz/s;

The *time-bandwidth product* $B\tau$ quantifies the pulse compression gain achievable through matched filtering (see Section **2.2.4**). For MARSIS $B\tau = 250$, which represents a 250× compression gain: as it well be shown in the following, the transmitted pulse (long, carries high energy), after matched filtering, is compressed to an effective duration of $\tau_c = 1\,\mu\text{s}$, reaching fine resolution capabilities.

2.2.4 Matched Filtering and Pulse Compression

The chirped waveform discussed in Section **2.2.3** achieves high energy and large bandwidth simultaneously, but the transmitted pulse remains long - 250 μs for MARSIS, corresponding to a spatial extent of 75 km. To realize the fine range resolution promised by the bandwidth ($\Delta\rho = c/2B = 150$ m), the received echo must be compressed through a processing technique that coherently sums energy across the entire bandwidth. This compression is achieved through *matched filtering*, an optimal linear filter that maximizes SNR while simultaneously achieving the theoretical range resolution limit.

When processing the received signal, a radar filters the signal prior to detection. A general filter can be described in terms of its impulse response, which is the output of the filter for an impulse input, or its frequency response, which is the Fourier transform of the impulse response. Thus, the spectrum of the receiver output y(t) is $Y(f) = H(f)R(f)$, where R(f) is the spectrum of the received signal. We want to maximize the SNR at a specific time T_M. The power of the signal component of the output at that instant is:

$$|y\left(T_M\right)^2| = \left|\int_{-\infty}^{\infty} S(f)H(f)e^{j2\pi ft}df\right|^2 \tag{2.50}$$

where S(f) is the spectrum of the transmitted waveform. Consider the interference as white noise with power spectral density $N_0/2$ watts per second. The total output noise power can be written as:

$$n_p = \frac{N_0}{2}\int_{-\infty}^{\infty} |H(f)|^2 df \tag{2.51}$$

The corresponding SNR at time T_M is:

$$\text{SNR} = \frac{|y\left(T_M\right)|^2}{n_p} = \frac{\left|\int_{-\infty}^{\infty} S(f)H(f)e^{j2\pi fT_M}df\right|^2}{\frac{N_0}{2}\int_{-\infty}^{\infty}|H(f)|^2 df} \tag{2.52}$$

The choice of H(f) that maximize the SNR can be determined using the Schwarz inequality:

$$\left|\int A(f)B(f)df\right|^2 \leq \left\{\int |A(f)|^2 df\right\}\left\{\int |B(f)|^2 df\right\} \tag{2.53}$$

the equality holds only when $B(f) = \alpha A^*(f)$, being α an arbitrary constant. Applying the inequality to the previous equation:

$$\text{SNR} \leq \frac{\int_{-\infty}^{\infty}|S(f)e^{j2\pi fT_M}|^2 df \int_{-\infty}^{\infty}|H(f)|^2 df}{\frac{N_0}{2}\int_{-\infty}^{\infty}|H(f)|^2 df} = \frac{2}{N_0}\int_{-\infty}^{\infty}|S(f)|^2 df \tag{2.54}$$

The SNR is maximized when:

$$H(f) = \alpha S^*(f)e^{-j2\pi fT_M} \quad \text{or} \tag{2.55}$$

$$h(t) = \alpha s^*(T_M - t) \tag{2.56}$$

Hence, the matched filter is defined as the linear filter whose impulse response h(t) is the time-reversed complex conjugate of the transmitted signal s(t). The waveform and the receiver filter needed to maximize the output SNR are a matched pair. The gain α is often set equal to unity, since it has no impact on the achievable SNR [2]. The time T_M at which the SNR is maximized is arbitrary.

The output of the matched filter is the cross-correlation between the received signal x(t) and the transmitted waveform:

$$y(t) = r(t) * h(t) = \int_{-\infty}^{\infty} r(\tau)s^*(T_M - \tau)d\tau \tag{2.57}$$

where $*$ denotes convolution. This integral computes the similarity between the received echo and a time-shifted replica of the transmitted signal at each time t. When r(t) contains a delayed and scaled version of s(t), as occurs when the echo returns from a point target, the matched filter output exhibits a sharp peak at the delay corresponding to the target range.

For a chirp waveform, the matched filter coherently integrates energy across the entire bandwidth B, compressing the long transmitted pulse into a much shorter output pulse whose duration is approximately $1/B$. This compression factor, equal to the time-bandwidth product $B\tau$, represents the pulse compression gain, the factor by which SNR is improved through matched filtering. Recalling Equation **2.46**, the matched filter impulse response is:

$$h(t) = s^*(-t) = \text{rect}\left(-\frac{t}{\tau}\right) \cdot \exp\left(-j2\pi f_0(-t) - j\pi\alpha(-t)^2\right) = \text{rect}\left(\frac{t}{\tau}\right) \cdot \exp\left(j2\pi f_0 t - j\pi\alpha t^2\right) \tag{2.58}$$

When the received echo $x(t) = A \cdot s(t - \tau_0)$ is convolved with $h(t)$, the output is the autocorrelation function R_{ss} of the chirp, centered at delay τ_0:

$$y(t) = A \cdot R_{ss}(t - \tau_0) \tag{2.59}$$

For a rectangular-envelope linear chirp as shown in Figure **2.6**, the autocorrelation function is approximately a sinc function:

$$R_{ss} \approx \tau \cdot \text{sinc}\left(Bt\right) \tag{2.60}$$

The matched filter compresses the pulse from length τ to a shorter effective length of $\tau_p = 1/B$. For this reason, matched-filter processing of the pulse signal is often termed *pulse compression* or range compression. The compression ratio - the factor by which pulse duration is reduced after matched filtering - is equal to the time-bandwidth product τB. The time-bandwidth product also determines the Doppler tolerance of the matched filter. For targets with significant radial velocity, the Doppler shift f_D introduces a mismatch between the received echo and the matched filter reference, degrading compression efficiency. The fractional SNR loss is negligible when $|f_D| \cdot \tau \ll 1$, imposing an upper limit on the pulse duration for high-velocity scenarios. For orbital radar sounders with relatively stable platform motion, Doppler shifts are typically small ($|f_D| \sim$ tens of Hz), and this constraint is easily satisfied.

The sinc-shaped autocorrelation function of a rectangular-envelope chirp exhibits range sidelobes, that are oscillations adjacent to the mainlobe peak. The first sidelobe is located at $1.5/B$ from the peak and has an amplitude of approximately ~ 13.2 dB relative to the mainlobe. Weighting functions suppress sidelobes by tapering the transmitted or received signal amplitude, trading mainlobe width (range resolution) for reduced sidelobe levels. The same window functions discussed in Section **2.1.4** for spectral analysis - Hamming, Taylor, Kaiser, and others - are employed here for the matched filter processing. The trade-off between sidelobe suppression and mainlobe broadening discussed previously applies equally to pulse compression: a Hamming window achieves ~ 43 dB sidelobes at the cost of $1.3\times$ mainlobe widening, while Taylor and Kaiser windows offer flexible design control over this balance [42].

While matched filtering is defined as a time-domain convolution, practical radar processors exploit the convolution theorem to implement it efficiently in the frequency domain, since convolution in time becomes multiplication in frequency. An alternative technique, particularly suited to hardware implementation, is *dechirping* or stretch processing (see next Section). The received chirp is mixed with a replica of the transmitted chirp (or its complex conjugate), producing a

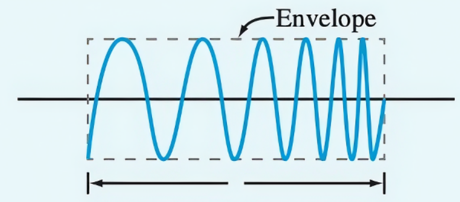

(a) Transmitted waveform

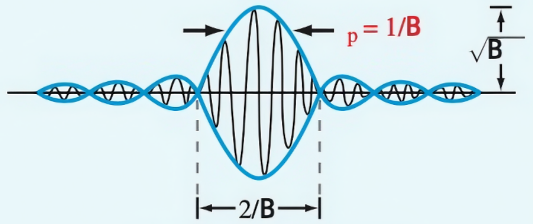

(b) Dechirped waveform

Figure 2.6: Waveforms for an LFM chirp radar with a point target. Note that the transmitted waveform has a boxcar envelope while the dechirped waveform has a sinc-function envelope. From [4].

constant-frequency beat signal whose frequency is proportional to target range. An FFT of this dechirped signal directly yields the range profile, combining mixing and matched filtering in a single operation. MARSIS employs this approach in its onboard analog front end, reducing the sampling rate required for the ADC [39].

2.2.5 Range Processing for Pulsed Chirp Signals

In the previous Section we introduced the concept of matched filtering, which consists of the compression of chirp pulses through cross-correlation in the frequency domain with the transmitted waveform. Even if this approach is theoretically easy to implement, in practice matched filtering on raw received data requires sampling the full instantaneous bandwidth at rates exceeding the Nyquist criterion, followed by computationally intensive processing on large data volumes. To overcome these issues, here we introduce the dechirping (or strech processing) procedure, which takes advantage of the chirp linear frequency modulation to reduce bandwidth and sampling rate requirements [49]. This technique is often applied to spaceborne systems where power, data rate, and onboard processing constraints are severe.

Dechirping operates by mixing the received chirp with a locally generated reference chirp identical (or conjugate) to the transmitted waveform. Consider again the expression for a transmitted-up chirp (Equation **2.46**), written in the exponential form:

$$s(t) = A \cdot \mathrm{rect}\left(\frac{t}{\tau}\right) \cdot \exp\left(j\,2\pi f_0 t + j\,\pi \alpha t^2\right) \tag{2.61}$$

and the received echo from a target at range R:

$$r(t) = A_r \cdot \mathrm{rect}\left(\frac{t-\tau_0}{\tau}\right) \cdot \exp\left(j\,2\pi f_0 (t-\tau_0) + j\,\pi \alpha (t-\tau_0)^2\right) \tag{2.62}$$

The reference chirp, generated by the receiver, is:

$$s_{\mathrm{ref}}(t) = \exp\left(-2\pi f_0 t - j\pi \alpha t^2\right) \tag{2.63}$$

which is the complex conjugate of the transmitted chirp. Mixing the received signal with the reference leads to:

$$r_{\text{dechirp}}(t) = r(t) \cdot s_{\text{ref}}(t) \tag{2.64}$$

Expanding and simplifying the exponential terms:

$$r_{\text{dechirp}}(t) = A_r \cdot \text{rect}\left(\frac{t-\tau_0}{\tau}\right) \cdot \exp\left(j\,2\pi f_0(t-\tau_0) + j\,\pi\alpha(t-\tau_0)^2\right) \cdot \exp\bigl(-j\,2\pi f_0 t - j\,\pi\alpha t^2\bigr) \tag{2.65}$$

$$= A_r \cdot \text{rect}\left(\frac{t-\tau_0}{\tau}\right) \cdot \exp\left(-j\,2\pi f_0\tau_0 + j\,\pi\alpha[(t-\tau_0)^2 - t^2]\right) \tag{2.66}$$

$$= A_r \cdot \text{rect}\left(\frac{t-\tau_0}{\tau}\right) \cdot \exp\left(-j\,2\pi f_0\tau_0 + j\,\pi\alpha\tau_0^2 - j\,2\pi\alpha\tau_0 t\right) \tag{2.67}$$

Consider the argument of the exponential term. The first two phase terms are constants (independent of t) and represent a fixed phase offset determined by the target range. The term $-j\,2\pi\alpha\tau_0 t$ corresponds to a constant beat frequency:

$$f_{\text{beat}} = \alpha\,\tau_0 = \frac{B}{\tau}\frac{2R}{c} \tag{2.68}$$

This formula shows that, after dechirping, each target at distance R produces a single-frequency sinusoid whose frequency is directly proportional to range. The chirp's quadratic phase has been removed, leaving only a linear phase term [2].

Let's now consider the bandwdith of the dechirped signal. In case of a multiple targets distributed over a range window ΔR, the dechirped signal contains several sinusoids at frequencies spanning:

$$\Delta f_{\text{beat}} = \frac{B}{\tau}\frac{2\,\Delta R}{c} \tag{2.69}$$

If we consider a range window of $\Delta R = 500$ km, $B = 1$MHz and $\tau = 250$ s, we obtain $\Delta f_{\text{beat}} = 13.3$ kHz, which is much smaller than the transmitted bandwidth. Following the Nyquist sampling criterion (Section **2.1.2**), the dechirped signal requires sampling at only $f_s > 2 \times 13.3 \approx 27$ kHz, which correpond to a reduction factor of nearly $40\times$ with respect to sampling the raw received chirp. This bandwidth reduction is the main advantage of strech processing: analog preprocessing compresses the signal bandwidth before digitization, dramatically relaxing ADC requirements [49].

It should be clear now that matched filtering and dechirping algorithms reach the same result through different processing architectures. Specifically, matched filtering is implemented in the following way:

- the received signal is sampled at $f_s > 2B$;
- compute FFT of the received signal: $R(f) = \mathcal{F}\{r(t)\}$;
- the signal is multiplied by the matched filter response: $Y(f) = R(f) \cdot H(f)$, with $H(f) = S^*(f)$;
- compute the inverse FFT: $y(t) = \mathcal{F}^{-1}\{Y(f)\}$.

This approach requires a high-speed ADC capable of sampling the full bandwidth B, and stores large data volumes for subsequent processing.

On the contrary, stretch processing achieves compression through a hybrid analog-digital approach:

- the signal is analog mixed with the reference chirp: $r_{\text{dechirp}}(t) = r(t) \cdot s_{\text{ref}}(t)$;
- a low-pass filter is applied to remove $2f_c$ components;
- the signal is sampled at $f_s > 2\,\Delta f_{\text{beat}} \ll 2B$;
- compute FFT to convert beat frequencies to range bins.

The final result is mathematically the same for both the approaches [50], the difference lies in when the compression is applied (before or after digitization). The equivalence can be understood by noting that multiplication by the reference chirp in the time domain (dechirping) followed by an FFT is equivalent to convolution with the matched filter in the time domain, or equivalently, multiplication in the frequency domain. Both operations correlate the received signal with the transmitted waveform; dechirping simply reorders the operations to exploit analog hardware for bandwidth reduction. Both approaches are widely used in modern radar systems. The choice between them mainly lies in the power and telemetry constraints of the instrument [43].

For moving platforms such as orbiting spacecraft, targets at different ranges may shift between range bins over the synthetic aperture integration time due to variations in slant range. This phenomenon, called *range cell migration* (RCM), becomes significant when the range change exceeds a fraction of the range resolution $\Delta\rho$. Range cell migration correction (RCMC) compensates for this effect, ensuring that energy from a given scatterer remains coherently integrated within a single range bin throughout SAR processing [35]. RCMC is typically implemented after range compression (i.e., after the FFT of the dechirped signal) and before azimuth compression. The correction is implemented by interpolating the range-compressed data to shift each target's energy back to its correct range bin as a function of azimuth position. For nadir-looking sounders operating at orbital altitudes, range migration is generally small over the integration time due to the nearly constant altitude and slow along-track motion. However, for high-resolution SAR imaging or when processing long synthetic apertures, RCMC becomes essential to avoid defocusing. The details of RCMC algorithms - including both time-domain and frequency-domain implementations - are discussed in Section **2.3.4** within the context of SAR processing.

2.3 Synthetic Aperture Radar (SAR) Principles

In the previous Section we showed the capability of a chirp radar to achieve fine range resolution through pulse compression. From Equation **2.42**, we demonstrated that range resolution is independent of the radar's distance from the target, depending only by the transmitted bandwidth. However, to generate high-resolution images in two and three dimensions from radar data, a target must be resolved also in the cross-range or *azimuth* direction, that is the direction perpendicular to the line-of-sight, typically aligned with the platform's velocity vector.

For conventional real aperture radar azimuth resolution is determined by the angular beamwidth of the antenna, which in turn depends on the physical size of the antenna aperture. As a consequence, fine azimuth resolution could be achieved only with impractically large antennas. The technique of *synthetic aperture radar* (SAR) [51] has been developed to address this issue by exploiting platform motion to synthesize an extremely long antenna aperture, achieving azimuth resolutions that are independent of range and orders of magnitude finer than real aperture systems.

2.3.1 The Azimuth Resolution Problem

Consider a side-looking radar mounted on a moving platform as shown in Figure **2.7**. The platform is at an altitude H above a flat surface and the radar illuminates it with a real aperture antenna of length L in the along-track direction. The angular beam width of the antenna can be written as:

$$\theta_{\text{az}} \approx \frac{\lambda}{L} \tag{2.70}$$

For a target at slant range R from the radar, the azimuth resolution, the minimum separation between two targets that can be distinguished in the azimuth direction, is determined by the footprint of this beam on the ground:

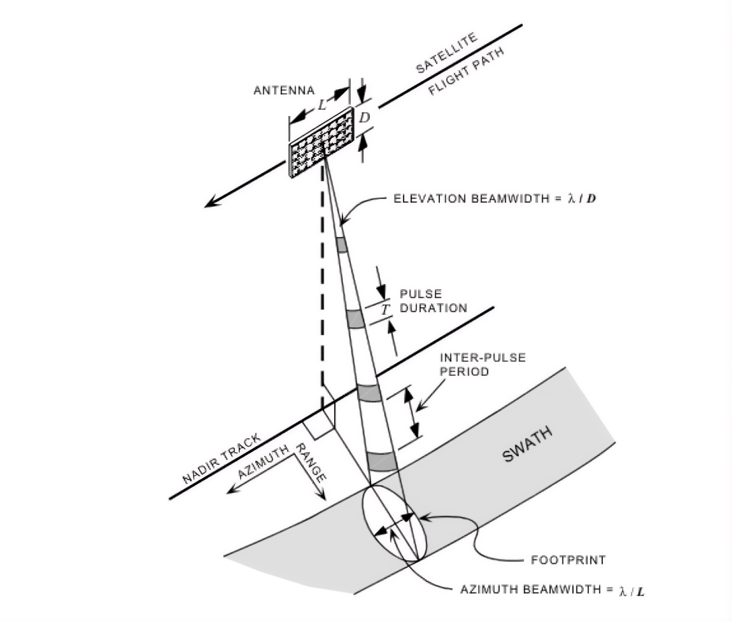

Figure 2.7: Observational geometry for a left looking SAR with a rectangular antenna. From [52].

$$\delta x_{\text{real}} = R \cdot \theta_{\text{az}} = \frac{\lambda R}{L} \tag{2.71}$$

It can be seen that azimuth resolution degrades linearly with range: increasing the observation wavelength, or building larger antennas are not feasible solutions in most of radar applications. The severity of this constraint becomes apparent when applied to orbital radar sounders. MARSIS operates at frequencies between 1.8 and 5 MHz, corresponding to wavelengths of $\lambda = c/f \approx 60-165$ m. At a typical orbital altitude of 300 km above the Martian surface, the range to the nadir surface is $R \approx 300$ km. The MARSIS dipole antennas have a total length of approximately $L = 40$ m. The resulting azimuth resolution is ≈ 280 km, considering the 4 MHz central frequency, while it degrades to over 600 km at the lowest MARSIS frequency of 1.8 MHz. To reach an azimuth resolution of $\delta x = 1$ km it would require an antenna length of more than 11 km. On the other hand, even at higher frequencies used by ice-sounding radars on Mars, like SHARAD at 20 MHz ($\lambda = 15$ m), achieving 100 m azimuth resolution from the same range would require a 2.25 km antenna. This example shows clearly the inadequacy of real aperture processing for radar planetary sounders and the need of syntethic aperture techniques.

2.3.2 The Synthetic Aperture Concept

Synthetic aperture radar is a technique that increase cross-range resolution without building a larger antenna. Instead, it exploits platform's motion to synthesize a very long antenna aperture. The target is sequentially illuminated from many different positions along the flight path. The echoes collected at each position are coherently combined in post-processing, effectively simulating the reception pattern of an antenna whose length equals the distance traveled during the observation [51]. Because data from multiple pulses are combined to form the effective high-resolution beam, the scene being imaged should not change while the data are collected, so that each pulse represents data from the same scenario. This again implies that SAR is intended primarily for imaging static scenes.

Figure **2.8** provides a visual explanation of this concept. Consider an array of N small antennas spaced along a line, each separated by distance Δx. If these antennas transmit and receive simultaneously, they form a real aperture of length $L_{\text{real}} = N \cdot \Delta x$, with azimuth resolution $\delta x \approx \lambda R / L_{\text{real}}$. Suppose then to replace this physical array with a single small antenna that moves along the same line, stopping at each of the N positions to transmit the signal and collect the echo. By coherently combining together all the echoes from all positions, we obtain the same result as using a simultaneous array. To perform the coherent sum, three requirements must be met:

- *phase stability*: the radar must maintain phase coherence across all pulse transmissions, requiring a stable oscillator and precise knowledge of platform motion;
- *motion compensation*: variations in platform trajectory must be measured and corrected in processing;
- *target stationarity*: the target must remain stationary relative to the surface during the integration time.

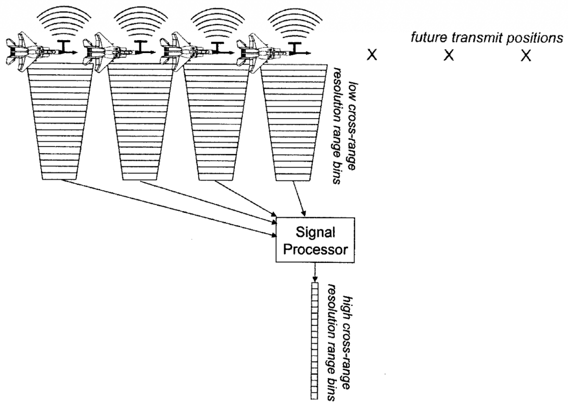

Figure 2.8: The concept of synthetic aperture operation. From [53].

Consider again Figure **2.7**, in which a radar is moving with velocity v along a straight line at constant altitude H over a point target. At time t, the along-track position of the radar relative to the target is $x = vt$, and the slant range to the target is:

$$R(t) = \sqrt{R_0^2 + (vt)^2} \tag{2.72}$$

Being R_0 the slant range to the target at its closest approach. The radial velocity is given by:

$$v_r(t) = \frac{dR}{dt} = \frac{v^2 t}{R(t)} \tag{2.73}$$

The instantaneous Doppler shift experienced by the target is:

$$f_D(t) = -\frac{2v_r(t)}{\lambda} = -\frac{2v^2 t}{\lambda R(t)} \tag{2.74}$$

where the factor of 2 accounts for two-way propagation. If the received signal is frequency analyzed at each time t, any energy observed in the return corresponding to range R and at Doppler frequency f_D, will be associated with a target at coordinate:

$$x_1 = \frac{\lambda\, R\, f_D}{2v} \tag{2.75}$$

while energy at different frequency f_D will be assigned to a corresponding coordinate x_2. Hence, even two targets are at the same range and in the same beam, they can be identified by analysis of the Doppler frequency spectrum of the received signal.

For a single target, Equation **2.74** describes its Doppler history during the entire observation time. Its Doppler frequency sweeps from positive to negative values in a continuous and predictable way ($f_D > 0$ for approaching target, $f_D < 0$ for receding one). For small angles around broadside, $R(t) \approx R_0$ and the Doppler varies approximately linearly with time:

$$f_D(t) \approx -\frac{2v^2}{\lambda R_0} t \tag{2.76}$$

Mathematically, this linear variation with time is analogous to the linear frequency modulation discussed in Section **2.2.3**. Just as range compression exploits the chirped structure of the transmitted pulse to achieve fine range resolution, azimuth compression exploits the Doppler chirp induced by platform motion to achieve fine azimuth resolution. The azimuth signal processing coherently integrates echoes across the synthetic aperture, compressing the initially broad antenna beam into a sharp response.

In analogy with range processing, the Doppler bandwidth determines the achievable azimuth resolution after processing. It is defined as the total range of Doppler frequencies experienced during synthetic aperture integration:

$$B_D \approx \frac{2v}{L} \tag{2.77}$$

This result can be derived by noting that the maximum Doppler occurs at the edge of the beam ($\theta = \pm\lambda/(2L)$), where the radial velocity is $v_r = v \sin\theta \approx v\,\lambda/(2L)$, giving $f_{D,max} = 2v_r/\lambda = v/L$. The full Doppler extent is thus $B_d = 2f_{D,\max}$ [35].

Thus, azimuth compression achieves the following resolution:

$$\delta x_{\mathrm{SAR}} = \frac{v}{B_D} = \frac{L}{2} \tag{2.78}$$

This counter-intuitive result states that improved resolution comes from smaller antennas, and it is completely independent of range, wavelength, and platform velocity.

2.3.3 SAR Geometry and Coordinate Systems

Synthetic aperture radar imaging requires careful consideration of the geometric relationship between the radar platform, the antenna beam, and the target surface. SAR systems like Earth-observing satellites or airborne platforms adopt a side-looking geometry: the radar antenna is oriented perpendicular to the flight direction and tilted to illuminate the surface at an oblique angle. In this scenario, SAR geometry is fully described by the following parameters:

- altitude H: the vertical distance from platform to surface;
- slant Range R: direct line-of-sight distance from radar to target;
- ground Range r_g: horizontal distance from nadir track to target;
- incidence angle θ_i: angle between the radar line-of-sight and the surface normal (vertical);
- look angle θ_L: angle between nadir and the radar line-of-sight (complementary to incidence angle for flat terrain: $\theta_L = 90° - \theta_i$.

For a flat surface, the relations between these quantities are:

$$R = \frac{H}{\cos\,\theta_i} = H\sqrt{1 + \left(\frac{r_g}{H}\right)^2} \tag{2.79}$$

$$r_g = H\,\tan\,\theta_i = \sqrt{R^2 - H^2} \tag{2.80}$$

and the slant range to nadir is $R_0 = H$. Side-looking geometry introduces several geometric distortion that must be taken into account. First of all, equal increments in slant range ΔR correspond to varying increments in ground range Δr_g:

$$\Delta r_g = \Delta R \cdot \sin\,\theta_i \tag{2.81}$$

this geometry distortion is called *range foreshortening*: steep incidence angles gives higher compression.

Layover effects occur when observing targets with significant vertical extent, such that the top and the bottom may have almost identical slant ranges, causing them to appear displaced or overlapped in the SAR image.

On the other hand, side-looking geometry provides also several advantages: the oblique incidence angle enhances sensitivity to surface roughness and slope, reduces nadir ambiguities, and separates surface returns from volume scattering.

Radar sounder or ground penetrating radar (GPR) like MARSIS or SHARAD adopt nadir-looking geometry: the surface is illuminated vertically beneath the spacecraft. This choice maximizes the signal penetration depth and avoid the geometric distortions previously described. The relationship between slant range and vertical depth depends on the dielectric properties of the subsurface. For a subsurface interface at depth d below the surface, the two-way travel time is:

$$\tau = \frac{2H}{c} + \frac{2d}{c/\sqrt{\epsilon_r}} \tag{2.82}$$

where d is the target depth and ϵ_r is the relative permittivity of the subsurface material. The apparent slant range is:

$$R_{\text{app}} = \frac{c\tau}{2} = H + d\sqrt{\epsilon_r} \tag{2.83}$$

Hence, conversion to true depth requires knowledge (or assumption) about subsurface permittivity.

The *instantaneous footprint*, the area on the surface illuminated by a single pulse, is determined by the antenna beamwidth and slant range. For a rectangular antenna of dimensions $L \times W$, the beamwidths in the along-track (azimuth) and cross-track (range) directions are approximately:

$$\theta_{\text{az}} \approx \frac{\lambda}{L}, \quad \theta_{\text{cross}} \approx \frac{\lambda}{W} \tag{2.84}$$

At slant range R, the footprint dimensions are:

$$F_{\text{az}} = R\,\theta_{\text{az}} = \frac{\lambda R}{L} \tag{2.85}$$

$$F_{\text{cross}} = \frac{R\,\theta_{\text{cross}}}{\sin\,\theta_i} = \frac{\lambda R}{W\sin\,\theta_i} \tag{2.86}$$

For nadir-looking geometry where $\theta_i \approx 0$, the cross-track footprint becomes very large, covering the entire width visible within the beamwidth. For side-looking SAR, the footprint is elongated in the cross-track direction (range direction) due to the oblique projection, with the elongation factor $1/\sin\,\theta_i$ increasing at shallow incidence angles.

As introduced in Section **1.5**, real antennas have a radiation pattern with a gradually decreasing gain as a function of the angle. For a uniformly illuminated aperture of length L, the one-way gain pattern is approximately:

$$G(\theta) = \text{sinc}^2\left(\frac{\pi\,L\sin\,\theta}{\lambda}\right) \tag{2.87}$$

This pattern has a narrow mainlobe of width $\theta_{\text{3dB}} \approx \lambda/L$ and sidelobes that decay with increasing angle. Scatterers near the edge of the beam contribute less to the received signal than those near boresight. This effect affects azimuth processing by creating a natural window function for azimuth compression. Also, the effective integration time is determined by the beamwidth: targets are coherently integrated only while within the main lobe. The antenna pattern also determines the ambiguous range intervals for pulsed systems: strong sidelobes can cause returns from distant ranges (outside the primary illuminated swath) to appear at incorrect ranges. Sidelobe suppression through antenna design and Doppler filtering is critical for reducing ambiguities in SAR images.

For nadir sounders, the broad, nearly omnidirectional beam pattern (necessary for low-frequency penetration) causes substantial overlap between successive footprints, providing the dense along-track sampling required for SAR processing but also contributing to surface clutter in subsurface returns.

It is worth noting that MARSIS and SHARAD employ dipole antennas, whose radiation pattern differs significantly from that of a rectangular aperture antenna described by Eq. **2.84**. Dipole antennas have a broad, nearly omnidirectional pattern in the plane perpendicular to the dipole axis, with half-power beamwidths of the order of tens of degrees. As a consequence, the instantaneous footprint of these instruments cannot be accurately described by Eq. **2.84**.

2.3.4 Range-Doppler Algorithm

The *range-Doppler algorithm* (RDA) is perhaps the most widely used image formation algorithm for processing SAR data [54]. The main steps of the RDA are the following:

- range compress the transmitted chirp to achieve range resolution $\Delta\rho = c/2B$;
- conversion to range-Doppler domain by FFT transforming the range-azimuth data;
- apply range migration correction to align targes at constant range;
- compress the azimuth response by azimuth matched filtering;
- transform back to range-azimuth domain to obtain the final focused image.

Each step is performed on the entire data matrix, making the algorithm computationally intensive but highly efficient due to FFT-based operations.

Usually, SAR data are acquired in two dimensions: fast time τ range, which refers the time scale between a single pulse, and slow time t, which refers to the time scale between consecutive pulses (determined by the PRF). Range compression has been described in Section **2.2.4**. The range-compressed signal from a point target at range R_0 and azimuth position x_0 can be written as:

$$s(\tau, t) = A \cdot \mathrm{sinc}\left[B\left(\tau - \frac{2(R(t)}{c}\right)\right] \cdot \exp\left[-j\frac{4\pi}{\lambda}R(t)\right] \tag{2.88}$$

Where $R(t) = \sqrt{R_0^2 + v^2t^2}$ is the instantaneous slant range and the sinc function represents the range compressed pulse. The azimuth position coordinate can be expressed as $x = vt$, making the data a function of $s(\tau, x)$.

RDA assumption is to threat range and azimuth domain data separately as one-dimensional processes. By applying an azimuth FFT data are converted to the range-Doppler domain:

$$S(\tau, f_D) = \int_{-\infty}^{\infty} s(\tau, t)\, \exp\left(-j\, 2\pi f_D\, t\right) dt \tag{2.89}$$

Or in the discrete form, with N azimuth samples:

$$S[\tau, k] = \sum_{n=0}^{N-1} s[\tau, n]\, \exp\left(-j\frac{2\pi\, k\, n}{N}\right) \tag{2.90}$$

where k is the Doppler frequency bin index, corresponding to the physical Doppler frequency $f_D = k \cdot f_s/N$, with $f_s = \mathrm{PRF}$ being the azimuth sampling rate. At this stage, each row of the data matrix corresponds to a specific range bin, and each column corresponds to a specific Doppler frequency. A point target, which occupied many azimuth samples in the range-azimuth domain (due to the broad antenna beam), now occupies a concentrated region in Doppler frequency - its energy is focused around the Doppler frequencies experienced as it transited through the beam. Those frequencies vary according to Equation **2.76**. The Doppler bandwidth, described by Equation **2.77**, is independent of range. Thus, all targets, regardless their distance, occupy the same Doppler bandwidth and can be compressed using the same matched filter.

Since the platform moves with respect to the target, the signal from a point-target would follow a hyperbolic path in the range-Doppler domain which can cross several range cells. For each azimuth position $x \neq 0$ (nadir), the range increases according to:

$$R(x) = \sqrt{R_0^2 + x^2} \tag{2.91}$$

The shift of the target between several range bins is called range cell migration (RCM) or range walk. The amount of range migration is range-dependent and increases with R_0. At higher orbits or for targets farther from nadir, migration becomes more severe. Conversely, migration decreases for shorter wavelengths (higher frequencies) and larger antennas (narrower beams, shorter integration time).

Range cell migration correction (RCMC) compensates for the range walk by shifting each Doppler frequency component to its appropriate range position. By combining Equation **2.76** and Equation **2.91**, the range at which a target appears for each Doppler frequency can be written as:

$$R(f_D) = \frac{R_0}{\sqrt{1 - \left(\frac{\lambda f_D}{2v}\right)^2}} \approx R_0\left[1 + \frac{1}{2}\left(\frac{\lambda f_D}{2v}\right)^2\right] \tag{2.92}$$

The approximation is valid for small Doppler frequencies. The range migration as a function of Doppler is:

$$\Delta R(f_D) = R(f_D) - R_0 \approx \frac{\lambda^2 f_D^2 R_0}{8v^2} \tag{2.93}$$

RCMC applies a range-dependent shift to each Doppler frequency bin to align all components at range R_0. In the range-Doppler domain, for each Doppler bin, the data is shifted by $-\Delta R(f_D)$ in the range direction. Since range samples are discrete, this requires interpolation (typically sinc interpolation or more efficient approximations such as polynomial interpolation).

After RCMC, all energy from a given scatterer is concentrated at a single range bin across all Doppler frequencies, enabling effective azimuth compression. The azimuth matched filter in the range-Doppler domain has the form:

$$H_{\text{az}}(f_D) = \exp\left[j\frac{4\pi R_0}{\lambda}\left(\sqrt{1-\left(\frac{\lambda f_D}{2v}\right)^2} - 1\right)\right] \tag{2.94}$$

This filter compensates for the quadratic phase variation induced by the hyperbolic range history *R(t)*. The focused SAR image in the range-azimuth domain can be written as:

$$s_{\text{focused}}(\tau, x) = \mathcal{F}^{-1}\{S(\tau, f_D) \cdot H_{\text{az}}(f_D)\} \tag{2.95}$$

The azimuth impulse response after matched filtering is:

$$h_{\text{az}}(x) \approx \frac{L}{2} \cdot \text{sinc}\left(\frac{2x}{L}\right) \tag{2.96}$$

And the 3 dB width of the sinc function gives the azimuth resolution of Equation **2.78**.

2.3.5 Focused and Unfocused SAR

In Section **2.3.4** we described the fully focused SAR processing: all echoes from a scatterer across the entire synthetic aperture are coherently integrated with precise phase correction, achieving the theoretical azimuth resolution limit $\delta x = L/2$. In practice, this simplification is almost never true. For instance, by assuming the radar platform's position in the x coordinate is $x = vt$, we have presumed that the radar platform is flying a straight, level, and constant velocity path. A full coherent SAR processing requires a precise knowledge of the orbit, the platform stability, and the processing complexity. For subsurface sounding applications, a simplified approach called unfocused SAR or partially focused SAR offers a practical compromise between resolution and computational cost. This technique performs coherent integration over a limited synthetic aperture length, shorter than the full beamwidth-limited aperture $L_{\text{syn}} = \lambda R/L$. This restriction gives several advantages:

- for shorter apertures, range migration $\Delta R \propto L_{\text{syn}}^2$ is minimal and may be neglected, eliminating the need for RCMC [35];
- without RCMC interpolation, processing reduces to range compression followed by simple coherent summation (or limited azimuth filtering) in the time domain;
- shorter integration times tolerate larger orbit errors, as phase errors accumulate over fewer pulses [55];
- avoiding RCMC and full matched filtering reduces processing time and memory requirements.

The downgrade is a coarser azimuth resolution. The azimuth resolution for unfocused processing with integration length L_{unfoc} is determined by the angular extent subtended by the synthetic aperture [35]:

$$\delta x_{\text{unfoc}} \approx \frac{\lambda R}{L_{\text{unfoc}}} \tag{2.97}$$

The output of unfocused (or focused) SAR processing before multi-looking or incoherent averaging is termed *single-look complex* (SLC) data. SLC data preserves the full complex-valued (amplitude and phase) information for each resolution cell [56], enabling interferometric processing, polarimetric analysis, and coherent change detection.

Multi-looking averages multiple adjacent SLC pixels (either in range, azimuth, or both) to reduce speckle noise at the cost of resolution [57]. For example, averaging N independent looks reduces resolution by $N\times$ but improves radiometric accuracy by $\sqrt{N}\times$. Multi-looked data is often converted to detected intensity $|s|^2$ amplitude $|s|$, discarding phase information.

To sum up, to perform full focused SAR processing the following requirements must be met:

- *precise Orbit Knowledge*. The azimuth matched filter assumes exact knowledge of the platform trajectory to compute the phase history. Orbit errors ΔR introduce phase errors:

  $$\Delta\phi = \frac{4\pi\Delta R}{\lambda} \tag{2.98}$$

  Specifically, phase errors must remain below $\pi/4$ radians to avoid significant SNR loss [58]:

  $$\Delta R < \frac{\lambda}{16} \tag{2.99}$$

- *Motion Compensation*. Beyond orbit errors, the platform may deviate from its trajectory due to attitude variations, or gravitational perturbations. These deviations must be measured and compensated in processing. Autofocus algorithms estimate residual phase errors directly from the SAR data by maximizing image sharpness or phase gradient metrics, providing an alternative when precise navigation data is unavailable [59]. Some of the most popular techniques are phase gradient autofocus (PGA), which estimates phase errors from main scatterers, and map drift (MD), which aligns sub-aperture images to estimate motion errors.

- *Phase Stability*.The radar's oscillator must maintain phase coherence over the integration time τ_{int}. Phase noise or frequency drift decorrelates the signal, degrading compression efficiency. The phase stability requirement is typically expressed as:

  $$\sigma_\phi < \frac{\pi}{4} \tag{2.100}$$

  where σ_ϕ is the RMS phase error over the synthetic aperture. Ultra-stable oscillators (USO), such as the one aboard Mars Express, provide frequency stability better than 10^{-11} over timescales of seconds to minutes, ensuring coherence over the full synthetic aperture for orbital radar sounders [39].

- *Target Stationarity*. The target must remain stationary relative to the surface during integration. For MARSIS, this requirement is easily met. In case of dynamic surfaces such as ocean waves, temporal decorrelation occurs over timescales of milliseconds to seconds, limiting the achievable synthetic aperture length and hence resolution. The decorrelation time depends on target type and environmental conditions [60].

Table **2.3** shows the trade off between the focused and the unfocused approaches.

Parameter	**Unfocused SAR**	**Focused SAR**
Integration length	$L_{\text{unfoc}} \ll L_{\text{syn}}$	$L_{\text{syn}} = \dfrac{\lambda R}{L}$
Azimuth resolution	$\delta x \approx \dfrac{\lambda R}{L_{\text{unfoc}}}$	$\delta x = \dfrac{L}{2}$
Range migration	Negligible (< 1 cell)	Significant (5–20 cells)
RCMC required	No	Yes
Orbit accuracy needed	10–50 m	< 5 m
Computational cost	Low	High
Typical applications	Subsurface profiling, reconnaissance	Surface imaging, interferometry

Table 2.3: Comparison between unfocused and focused SAR processing

2.4 Radar Focusing and Backprojection

In Section **2.2** and Section **2.3** we established the principles of pulse compression and synthetic aperture radar, demonstrating how matched filtering achieves fine resolution in both range and azimuth dimensions. In this Section we describe the main focusing algorithms used in radar subsurface imaging, ranging from frequency-domain methods based on simplified geometric assumptions, such as Stolt migration and phase-shift, to time-domain backprojection, which we implemented in the frequency domain to elaborate 3D focused MARSIS images during this project.

Consider a point target illuminated by a GPR, like in Figure **2.9**. In an unfocused radar image, the echo received doesn't appear as a single point, but as a hyperbolic signature, stretched across multiple range-azimuth pixels. This effect occurs because the transmitting antenna illuminates the scatterer not only when it passes exactly over it, but also at other positions, as long as the target is inside the main antenna beam. For each of these positions, the antenna receives an echo at a different range and azimuth location. The collection of these echoes traces out a characteristic hyperbolic curve in the range-azimuth plane (Section **2.3.4**) [61].

Focusing, also called migration in the geophysical framework, is the process of collapsing these hyperbolic signatures back to their true point locations, reconstructing the spatial distribution of scatterers from the distorted radar measurements [62]. For subsurface radar sounders operating in complex three-dimensional environments, focusing is essential for accurate geolocation, resolution enhancement, and interpretation of geological structures. To accurately reconstruct the image with a focusing algorithm, the radar data must satisfy spatial sampling criteria in the azimuth direction. The required platform position spacing $\Delta\epsilon$ is determined by the Nyquist criterion (Equation **2.14**) applied to the spatial frequencies present in the focused image. The maximum spatial frequency in the image is determined by the radar wavelength and the maximum scattering angle. For backscatter (monostatic) geometries, the maximum wavenumber is approximately [4]:

$$k_{\max} = \frac{4\pi}{\lambda} \tag{2.101}$$

To sample correctly the signal, Nyquist criterion requires at least 2 samples per wavelength in the spatial domain:

$$\Delta\epsilon < \frac{\lambda}{2} \tag{2.102}$$

Oversampling is common practice to avoid aliasing artifacts near the Nyquist limit and to accommodate velocity variations.

Before proceeding further, it is important to distinguish focusing from tomographic reconstruction. Both are inverse problems, but they differ in the number of datasets required and the type of ambiguity they address: focusing operates on a single dataset and addresses the within-plane ambiguity by collapsing diffraction signatures to their true spatial locations, while tomography combines multiple datasets acquired from different viewing angles to resolve the cross-plane ambiguity and reconstruct the reflectivity distribution in 2D or 3D. The backprojection algorithm appears in both contexts, but the underlying inverse problems are distinct, as summarized in Table **2.4**.

Aspect	**Focusing (Migration)**	**Tomographic Reconstruction**
Problem scope	Single-dataset problem	Multi-dataset problem
Input data	Unfocused radar data from one orbit pass	Multiple radar datasets from different viewing angles
Objective	Collapse diffraction signatures to their true spatial locations	Reconstruct reflectivity distribution from angular diversity
Output	Focused image with improved resolution and localization	2D or 3D reflectivity distribution
Problem type	Linear inverse problem, single viewpoint	Linear inverse problem, multiple viewpoints
Regularization required	Yes	Yes
Main constraints	Angular coverage limited to single orbit	Angular coverage, data consistency, noise sensitivity

Table 2.4: Comparison between focusing (migration) and tomographic reconstruction in radar imaging

2.4.1 Hyperbolic Signatures from Point Scatterers

Consider a radar platform moving with velocity v along a straight trajectory at altitude H above a point scatterer located at position (x_0, y_0, z_0) in a Cartesian coordinate system where x is along-track (azimuth), y is cross-track, and z is vertical (depth positive downward). The radar position at slow time t is $(vt, 0, 0)$ for a nadir-looking geometry. The range from the radar to the scatterer is:

$$R(t) = \sqrt{(vt - x_0)^2 + y_0^2 + (H + z_0)^2} \tag{2.103}$$

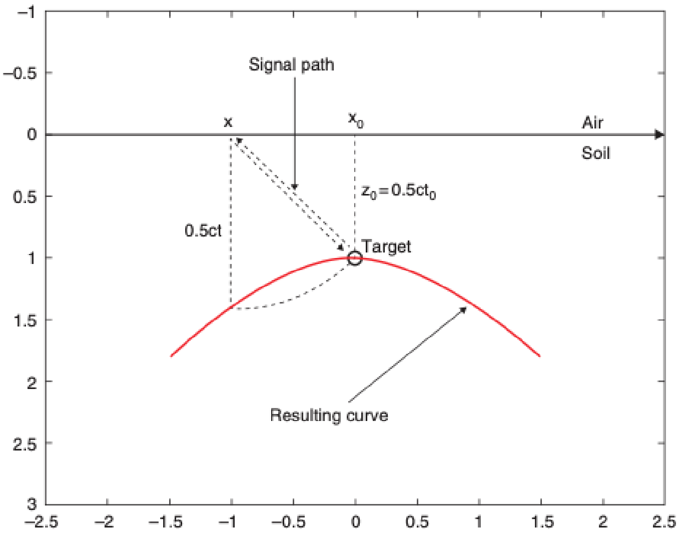

Figure 2.9: Mechanism of the radar observation and the forming principle of diffraction hyperbolas. From [63].

With the corresponding echo received after the two-way travel time $\tau(t) = 2R(t)/c$. In the range-azimuth radar image (after range compression but before azimuth focusing), the scatterer's energy is distributed along the locus of points satisfying this range-time relationship. For a scatterer directly beneath the flight track ($y_0 = 0$) the relationship simplifies to:

$$\tau(t) = \frac{2}{c}\sqrt{(vt - x_0)^2 + (H + z_0)^2} \tag{2.104}$$

That is the equation of a hyperbola in the (x, τ) plane, where $x = vt$ is the along track position. The hyperbola has a its minimum at $x = x_0$, where:

$$\tau_{\min} \approx \frac{2(H + z_0)}{c} \tag{2.105}$$

and the asymptotes are:

$$\frac{d\tau}{dx} \approx \pm\frac{2}{c} \tag{2.106}$$

where the factor $2/c$ accounts for the two-way travel time. This asymptotic approximation holds in the limit $|vt - x_0| \gg (H + z_0)$, i.e., when the along-track distance from the scatterer is much larger than the altitude, so that the hyperbola approaches its asymptotes. The formula holds also for $y_0 \neq 0$, and the curvature of the hyperbola depends on the depth of the scatterer: deeper scatterers produce less curved hyperbolae (approaching straight lines at infinite depth), while shallow scatterers produce sharply curved hyperbolae. The presence of these diffraction patterns complicates deeply the interpretation of GPR images: in practice, when multiple scatterers are present, their hyperbolae overlap, creating complex interference patterns that obscure the true spatial distribution of scatterers. In case of extended planar interfaces, hyperbolic signatures appear only at discontinuities such as edges, faults, or roughness features. Continuous planar interfaces parallel to the surface do not produce hyperbolae; they appear as simple range-delay lines. However, any lateral variation, such as buried channel or a dipping layer, generates diffraction hyperbolae that must be focused to reveal the true geometry [64].

SAR processing methods and, specifically, the RDA algorithm shown in Section **2.3.4**, perform azimuth compression assuming that range variations are solely due to the synthetic aperture geometry, not due to scatterer depth or off-nadir position [35]. The major assumption of these methods is that targets lie on a planar reference surface (typically the mean surface elevation), which is not valid in subsurface radar sounding scenario. A buried reflector experiences a different range history than a surface reflector at the same horizontal position. The implementation of the RDA's azimuth matched filter, designed for the surface geometry, mismatches the subsurface geometry, causing incomplete focusing (subsurface scatterers remain smeared), geometric distortion (scatterers are placed in incorrect positions), and amplitude reduction (echoes from the same subsurface target are not coherently summed).

Given this context, the focusing problem can be stated as follows: *given unfocused radar measurements $s(\tau, x)$ containing hyperbolic diffraction signatures, determine the spatial distribution of scatterers $\rho(x, z)$ that produced those measurements* [65]. From a mathematical point of view, consider several point scatterers with reflectivities ρ_i at positions (x_i, z_i). The unfocused radar image is:

$$s(\tau, x) = \sum_i \rho_i \cdot h\left(\tau - \frac{2R_i(x)}{c}\right) \tag{2.107}$$

where $h(\tau)$ is the compressed pulse shape and $R_i(x)$ is the range to the i-th scatterer, and the sum is over all scatterers. Focusing algorithms inverts this forward scattering relationship. For each scatterer, it computes the theoretical hyperbolic trajectory expected, sums the measured data along this trajectory, and assigns the summed value as the focused image intensity at (x_0, z_0). This summation process coherently integrates energy along the diffraction hyperbola, collapsing it to a point response when the test position coincides with a true scatterer [64].

2.4.2 Time-Domain Backprojection

Time-domain backprojection is one of the most versatile and geometrically intuitive methods for focusing radar data. Its theoretical foundation lies in Kirchhoff migration [66], which is based on the Kirchhoff integral theorem, a formal solution to the scalar wave equation governing EM wave propagation, physically equivalent to the Huygens principle: each point on a wavefront acts as a secondary source of spherical waves, and the total field is reconstructed by coherently summing all such contributions. For a wave function in a homogeneous medium of velocity v, the Kirchhoff integral theorem gives the reconstructed time-dependent wavefield at image point (x, z) as [66]:

$$P(x, z, t) = \frac{1}{2\pi} \int_{-\infty}^{\infty} \frac{\cos\theta}{vR} \cdot \frac{\partial}{\partial t} P\Big(x, z, t - \frac{r}{v}\Big)\, dx \tag{2.108}$$

where $r = \sqrt{(x - x_m)^2 + z^2}$ is the distance from the image point to the observation point $(x_m, 0)$, θ is the angle between the propagation direction and the depth axis, and R is the geometric spreading factor. Compared to pure hyperbolic summation, Kirchhoff migration additionally accounts for spherical spreading ($1/\sqrt{vr}$ for 2D, $1/(vr)$ for 3D), the directivity factor $\cos\theta$, and phase corrections, improving the amplitude fidelity of the focused image [66].

Time-domain backprojection is the practical discrete implementation of Eq. **2.108**. Whereas $P(x, z, t)$ reconstructs the full time-dependent wavefield, the backprojection image $I(x, z)$ is obtained by evaluating the field at the correct propagation delay $\tau(x, z, \epsilon)$ for each platform position, effectively collapsing the time dimension into a focused spatial image. The correction factors of Kirchhoff migration - spherical spreading, directivity, and phase - are absorbed into a general weighting function w. Compared to frequency-domain methods, which require varying assumptions about squint, range curve migration, platform motion, and point spread homogeneity, backprojection makes fewer assumptions, as detailed below.

Consider unfocused radar data $s(\tau, \epsilon)$ where τ is the fast time and ϵ the slow time. The backprojection integral for forming a focused image $I(x, z)$ at spatial position (x, z) is:

$$I(x, z) = \int_{-\infty}^{\infty} s(\tau(x, z, \epsilon), \epsilon) \cdot w(x, z, \epsilon) d\epsilon \tag{2.109}$$

where $\tau(x, z, \epsilon)$ is the two-way travel time from the radar at position ϵ to the image point (x,z) and back, and $w(x, z, \epsilon)$ is a weighting function that accounts for geometric spreading, antenna pattern, obliquity factors, and other amplitude corrections, as stated before. The integral sums contributions from all platform positions ϵ along the synthetic aperture. The backprojection algorithm reverses the diffraction process seen in Section **2.4.1**: whereas a point scatterer at (x, z) creates energy distributed along the hyperbola $\tau(\epsilon)$, backprojection samples the measured data along this hyperbola and sums the values, concentrating energy back at the scatterer position.

To understand how this algorithm works, consider a scatterer at position (x_0, z_0) generating a scattered wavefield that propagates as a spherical wave. The radar, located at various positions ϵ along its trajectory, samples the wavefield at different times corresponding to the propagation delay $\tau(x_0, z_0, \epsilon)$. When the test position (x, z) coincides with the location (x_0, z_0) of a scatterer, the backprojection summation integrates coherently: all samples along the diffraction curve correspond to echoes from the same target, arriving in phase with each other. This coherent sum produces a larger amplitude in the focused image. When the test position does not coincide with a scatterer, the samples along the corresponding hyperbola come from different physical locations (or noise) and have random phase relationships. The contributions add destructively or incoherently, producing low amplitude that averages toward zero. Fundamentally, backprojection acts as a spatial matched filter, achieving maximum gain when the test position coincides with a true target location [67].

The weighting function $w(x, z, \epsilon)$ corrects for several physical effects:

- as shown in Section **1.5**, the received power decreases quadratically with range. To compensate this effect, the backprojection weight includes a range-dependent gain:

  $$w_{\text{geom}}(x, z, \epsilon) = R(x, z, \epsilon)^n \tag{2.110}$$

  where n depends on the dimensionality of spreading: $n = 2$ for 3D spherical spreading (monostatic surface radar), $n = 1$ for 2D cylindrical spreading (line sources), or more complex expressions for bistatic geometries [68]. For subsurface radar sounders, $n = 2$ is typical.

- According to the definition of antenna gain $G(\theta)$, the intensity of signal decreases at the edges of the antenna beam. Thus, the weighting function includes:

  $$w_{\text{ant}}(x, z, \epsilon) = G(\theta(x, z, \epsilon)) \tag{2.111}$$

  With this definition, the weighting term acts as a window function, reducing sidelobes in the focused image [61].

The full weighting function is then:

$$w(x, z, \epsilon) = w_{\text{geom}} \cdot w_{\text{obl}} \cdot w_{\text{ant}} \cdot w_{\text{other}} \tag{2.112}$$

where the term w_{obl} accounts for oblique geometry, while w_{other} can include additional corrections for attenuation, polarization effects, or system calibration factors.

In real applications radar signals acquired are discrete, in the form $s[\tau_i, \epsilon_j]$, where i considers range sample and j indexes azimuth positions. The continuous backprojection integral becomes a discrete summation:

$$I[m,n] = \sum_{j=1}^{N_\epsilon} s[\tau_{mn}(j), j] \cdot w[m,n,j] \tag{2.113}$$

where [m,n] indicate the output image pixel at position (x_m, z_n), and $\tau_{m,n}(j)$ is the travel time from platform position j to pixel $[m,n]$. Interpolation must be considered with caution: the computed travel time $\tau_{m,n}(j)$ does not necessarily coincide with a discrete sample τ_i. In this case, the data value $s[\tau_{mn}(j), j]$ must be obtained by interpolating between adjacent samples. The most used interpolation methods are nearest neighbor ($s[\tau, j] \approx s[i_{\text{nearest}}, j]$), linear ($s[\tau, j] = (1-\alpha)s[i,j] + \alpha s[i+1, j]$), and sinc ($s[\tau, j] = \sum_k s[k,j] \cdot \text{sinc}\left[(\tau - \tau_k)/\Delta\tau\right]$) interpolation. For radar processing, linear interpolation typically provides sufficient accuracy with reasonable computational cost [69].

Time-domain backprojection possesses several advantages with respect to the frequency-domain methods:

- backprojection utilizes the same algorithm regardless of imaging mode. Frequency-domain methods, on the other hand, typically require modification depending on the imaging mode and geometry.
- By explicitly incorporating the imaging geometry, backprojection implements a pixel-dependent matched filtering process that closely models the expected radar return. Higher-order effects, such as pulse-by-pulse antenna gain variations in phase and magnitude, are therefore handled naturally. Frequency-domain methods, being convolution-based, apply a uniform response across the scene and can only approximately compensate for these effects.
- Motion compensation in the backprojection algorithm is implicit, since it utilizes available ground information (see the previous point). This advantage becomes more significant at lower airborne altitudes.
- Imaging geometries involving widely separated transmit and receive platforms with non-parallel motion are generally incompatible with conventional methods, whereas backprojection can accommodate them straightforwardly when platform positions are known.
- Backprojection imposes no fundamental constraints on the imaging swath beyond those imposed by the radar geometry and underlying physical limits, such as minimum PRF requirements. In particular, it naturally accommodates imaging configurations in which the relative velocity between the antenna and the scene varies across the swath. Frequency-domain methods, by contrast, are commonly restricted in this respect.
- Because frequency-domain methods rely on transformations to the frequency-domain, they inherently require buffering large data blocks, with memory usage scaling with synthetic aperture length. Backprojection avoids this requirement by processing echoes sequentially, enabling pulse-by-pulse imaging without storing the full raw dataset.
- Backprojection enables efficient reconstruction of selected regions of interest, with computational cost scaling directly with the number of imaged pixels, whereas frequency-domain approaches generally require full-scene processing or provide only marginal savings when imaging a subset.

While advances in computational hardware have largely mitigated the computational burden of time-domain backprojection, its main limitation remains the requirement for precise knowledge of the imaging geometry, including accurate antenna trajectories and terrain elevation models. In contrast, frequency-domain methods rely on simplified geometric assumptions and operate in range coordinates defined by the radar measurements; although computationally efficient, their robustness decreases when these assumptions are violated.

2.4.3 Frequency-Domain Migration Methods

Frequency-domain migration methods take advantage of the FFT algorithms to solve the focusing problem as a coordinate transformation or phase correction in the spectral domain [62]. While more efficient from a computational point of view, frequency-domain methods require severe assumption about imaging geometry, velocity structure, and target distribution. In this Section we provide a description of the three most canonical frequency-domain migration approaches: Stolt migration, shift migration, and omega-k algorithm.

Stolt Migration (f-k Migration)

Stolt migration, also termed frequency-wavenumber migration or f-k migration, is the classic frequency-domain method for focusing radar data [62]. The key idea is interpreting migration as a coordinate transformation in the 2D frequency-wavenumber domain, converting the data from measurement coordinates (temporal frequency ω, spatial wavenumber k_x) to image coordinates (spatial wavenumbers k_x, k_z).

Stolt migration is based on the exploding reflector model (Figure **2.10**), which consists in assuming that all the reflectors in the medium explode simultaneously at $t = 0$ and become upward-emitting sources. This postulate reduces the back and forth wave propagation to an upward wave scenario. [70]. Under this assumption, migration is reduced to reverse wave propagation and determining the wavefield at $t = 0$, the time of the explosion, at the scatterer location [71].

Consider a constant velocity medium with wave speed c. The wavefield $\Psi(x, z, t)$ satisfies the scalar wave equation:

$$\frac{\partial^2 \Psi}{\partial x^2} + \frac{\partial^2 \Psi}{\partial z^2} = \frac{1}{c^2} \frac{\partial^2 \Psi}{\partial t^2} \tag{2.114}$$

By applying the Fourier transform $\tilde{\Psi}(k_x, z, \omega)$ with respect to x and t, the wave equation becomes:

$$\frac{\partial \tilde{\Psi}}{\partial z^2} + k_z^2 \tilde{\Psi} = 0 \tag{2.115}$$

Taking the reference surface at $z = 0$, the solution to the transformed wave equation is:

$$\tilde{\Psi}(k_x, z, \omega) = \tilde{\Psi}(k_x, 0, \omega) \cdot e^{-jk_z z} \tag{2.116}$$

At the imaging time $t = 0$ (*zero-offset imaging condition*), the migrated image is obtained by taking the inverse Fourier transform:

$$I(x, z) = \mathcal{F}^{-1}_{k_x, k_z} \{\tilde{\Psi}(k_x, 0, \omega(k_z))\} \tag{2.117}$$

Stolt migration takes advantage of a coordinate transformation to map the data in the image space (Stolt mapping):

$$k_z = \sqrt{\left(\frac{2\omega}{c}\right)^2 - k_x^2} \tag{2.118}$$

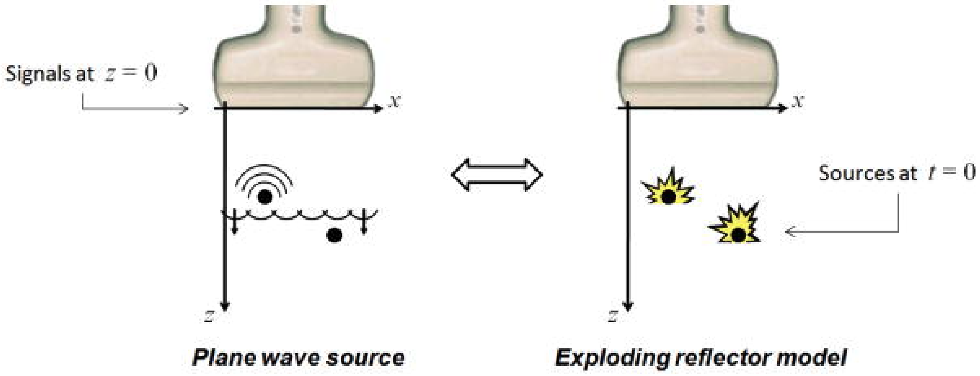

Figure 2.10: The exploding reflector model (ERM). From [70].

this dispersion relation is the basis of all frequency-domain algorithm, which relates the temporal frequency ω with the horizontal and vertical wavenumbers k_x, k_z.
In the end, the Stolt migration algorithm workflow can be summarized as follows:

- transform the input data $s(\tau, x) \rightarrow S(\omega, k_x)$ by applying a 2D FFT;
- resample $S(\omega, k_x)$ onto the new grid $S(k_x, k_z)$ through Equation **2.118**. This step requires 2D interpolation in the spectral domain;
- apply the inverse 2D FFT $S(k_x, k_z) \rightarrow I(x, z)$ to obtain the final focused image in spatial coordinates.

Stolt migration requires several severe assumption to be valid, that limits its range of applications to specific scenarios. First of all, it is supposed that the entire subsurface of interest has uniform wave speed. Velocity variations (corresponding to change in permittivity) cause phase errors and defocusing. Equation **2.118** holds only for radar monostatic configurations. Finally, the reference surface is assumed to be flat, since topographic variations change the traveltime relationship, causing mispositioned reflectors. In MARSIS framework, Stolt migration offers low flexibility and applicability, due to the presence of a stratified subsurface with different permittivity values and topographic variations. On Earth, this algorithm has been extensively used for ice-sounding radars in Antarctica, where all the previous requirements are easily met. Furthermore, as all frequency-domain migration methods, Stolt migration operates under the Born approximation (Section **2.5.2**), which limits the reconstruction to a qualitative estimate of the reflectivity distribution, without recovering quantitative permittivity values.

Phase-Shift Migration

Phase-Shift Migration [65] can be seen as an improved version of Stolt migration, since it takes into account for laterally varying velocity $c(x)$. Phase-shift migration proceeds iteratively in depth:

- apply a 2D FFT to the data: $s(\tau, x) \rightarrow S(\omega, k_x)$;
- consider the depth axis z as the sum of several depth layers Δz, such that $z = 0, \Delta z, 2\Delta z, ..., z_{\max}$. For each depth layer find k_z using Equation **2.118**;
- the solution of the one-way equation for downward propagation in the frequency domain for each depth layer Δz is:

$$\tilde{\Psi}(k_x, z + \Delta z, \omega) = \tilde{\Psi}(k_x, z, \omega) \cdot e^{-jk_z \Delta z} \tag{2.119}$$

Where the exponential term is the phase-shift operator, which is applied to the wavefield;

- an image is extracted for each layer by applying an inverse FFT. The imaging condition $t = 0$ is met by summing over all frequency components of the extrapolated wavefield at each depth step. The final image $I(x, z)$ is formed by stacking together all the depth layer images.

In case of lateral velocity variations $c(x)$, the phase-shift operator must be applied separately for each x position with the local $k_z(x)$. This split-step approach handles lateral velocity variations in the medium, even if it increases the algorithm computational cost.
Despite being more flexible than Stolt migration, phase-shift migration is still difficult to apply to MARSIS data: the algorithm requires constant velocity within each depth layer Δz, and the flat surface assumption must be valid. Also, we recall that the Born approximation underlies the entire formulation, limiting the reconstruction to a qualitative imaging result.

Omega-k Algorithm (Range Migration Algorithm)

The *omega-k algorithm*, also called the range migration algorithm, is a variant of frequency-domain migration specifically optimized for SAR processing [72].
The algorithm outline is the following:

- FFT transform range-compressed data $s(\tau, x) \rightarrow S(\omega, k_x)$;
- apply the RCMC (see Section **2.3.4**) to shift each Doppler component to the correct range;
- apply azimuth matched filter (it can be done either in time and frequency domains);
- apply the 2D inverse FFT to obtain the final focused SAR image:

$$S(\omega, k_x) \rightarrow I(\tau, x) \tag{2.120}$$

In practice, omega-k algorithm is designed for side-looking SAR, and it has low applicability for subsurface sounding, since it assumes flat surfaces and a constant velocity layer.
Table **2.5** summarizes the main characteristics of all the imaging methods discussed in this Section, taking into account imaging geometry, velocity structure, and computational constraints. It is worth noting that all the described algorithms represent approximate solutions to the inverse scattering problem (see next Section): they differ in the type and degree of approximation introduced, ranging from the geometric simplifications of frequency-domain methods to the Born approximation underlying all of them, which limits the reconstruction to a qualitative estimate of the reflectivity distribution. A detailed analysis of these algorithms has been necessary for the development of this work, since the MARSIS framework is unique with respect to the contexts for which these methods were originally developed, and understanding which approach best fits our data has required a careful comparative evaluation.

2.5 Inverse Scattering Theory

The radar imaging and tomography algorithms described in the previous sections are practical implementations of solutions to an inverse scattering problem: given measurements of electromagnetic waves scattered by subsurface structures, reconstruct the spatial distribution of those structures. To understand the theoretical foundations, approximations, and limitations of these algorithms,

Criterion	**Time-Domain Backprojection**	**Stolt (f–k)**	**Phase-Shift**	ω–k
Computational cost	High ($\mathcal{O}(N^3)$)	Low ($\mathcal{O}(N^2 \log N)$)	Medium ($\mathcal{O}(N_z N^2 \log N)$)	Low ($\mathcal{O}(N^2 \log N)$)
Geometric flexibility	Excellent (arbitrary geometry)	Poor (flat geometry only)	Moderate (flat layered media)	Poor (flat geometry only)
Velocity structure	Arbitrary $c(x, z)$	Constant c only	Laterally varying $c(x)$	Constant c only
Topography handling	Native	Requires preprocessing	Requires preprocessing	Requires preprocessing
Bistatic geometry	Native	Not supported	Not supported	Not supported
Typical applications	GPR, irregular terrain, bistatic systems	Ice sheets, flat surfaces	Stratified media	Spaceborne SAR

Table 2.5: Comparison of wave-equation-based migration algorithms.

we must first establish the rigorous mathematical framework of scattering theory [73] [74]. This Section develops the forward scattering problem (relating object properties to scattered fields), formulates the inverse problem as a linear operator equation, analyzes its ill-posedness and regularization requirements, and connects this abstract framework to practical radar processing algorithms like backprojection and tomography.

2.5.1 The Forward Scattering Problem

The direct or *forward scattering problem* consists in the calculation of the field scattered from a known dielectric permittivity profile, under the radiation of electromagnetic waves radiated by a known sources [63]. The *inverse scattering problem* consists in the reconstruction given the knowledge of the scattered field.
Let's start from the mathematical description of the forward scattering problem. Consider a radar transmitting an electromagnetic wave into a medium containing subsurface scatterers. The total electric field is given by the vector sum of the incident and the scattering field, such that: $\boldsymbol{E}_{\text{total}}(\boldsymbol{r}) = \boldsymbol{E}^i(\boldsymbol{r}) + \boldsymbol{E}^s(\boldsymbol{r})$. $\boldsymbol{E}_{\text{total}}$ in a medium with spatially varying permittivity $\epsilon_r(\boldsymbol{r})$ satisfies the inhomogeneous Helmholtz equation:

$$\nabla^2 \boldsymbol{E}_{\text{total}} + k_b^2 \frac{\epsilon_r(\boldsymbol{r})}{\epsilon_{r,b}} \boldsymbol{E}_{\text{total}} = 0 \tag{2.121}$$

where $k_b = \omega\sqrt{\mu_0 \epsilon_0 \epsilon_{r,b}} = 2\pi/\lambda_b$ is the background wavenumber, $\lambda_b = c/(\sqrt{\epsilon_{r,b}} f)$ is the wavelength in the background medium, and $\epsilon_{r,b}$ is the background relative permittivity. We threat the scatterer as a background perturbation. By adding and subtracting $k_b \boldsymbol{E}_{\text{total}}$, we separate the two contributions:

$$\nabla^2 \boldsymbol{E}_{\text{total}} + k_b^2 \boldsymbol{E}_{\text{total}} + k_b^2 \left[\frac{\epsilon_r(\boldsymbol{r})}{\epsilon_{r,b}} - 1 \right] \boldsymbol{E}_{\text{total}} = 0 \tag{2.122}$$

$$\nabla^2 \boldsymbol{E}_{\text{total}} + k_b^2 \boldsymbol{E}_{\text{total}} = -k_b^2 \left[\frac{\epsilon_r(\boldsymbol{r})}{\epsilon_{r,b}} - 1 \right] \boldsymbol{E}_{\text{total}} \tag{2.123}$$

With the right-hand side acting as a source term for scattering. According to [73], we define the *scattering potential*, or object potential, as:

$$V(\boldsymbol{r}) = k_b^2 \frac{\Delta\epsilon_r(\boldsymbol{r})}{\epsilon_{r,b}} = k_b^2 \rho(\boldsymbol{r}) \tag{2.124}$$

with $\Delta\epsilon_r(\boldsymbol{r}) = \epsilon_r(\boldsymbol{r}) - \epsilon_{r,b}$, and $\rho(\boldsymbol{r})$ is the reflectivity, already defined in Section **1.3.1** as the normalized dielectric contrast. Thus, the wave equation can be rewritten in its canonical form as:

$$\nabla^2 \boldsymbol{E}_{\text{total}} + k_b^2 \boldsymbol{E}_{\text{total}} = -V(\boldsymbol{r})\boldsymbol{E}_{\text{total}} \tag{2.125}$$

It is clear from this formulation that estimating the scattering potential from radar measurements-which is the inverse problem statement - allows us to find the value of the reflectivity, which yields to the absolute permittivity of the scatterer:

$$\epsilon_r(\boldsymbol{r}) = \epsilon_{r,b}[1 + \rho(\boldsymbol{r})] \tag{2.126}$$

The exact solution for the total field in the presence of scatterers can be written as the Lippmann-Schwinger integral equation [75]:

$$\boldsymbol{E}_{\text{total}}(\boldsymbol{r}) = \boldsymbol{E}^i(\boldsymbol{r}) + \int_V G(\boldsymbol{r}, \boldsymbol{r'})V(\boldsymbol{r'})\boldsymbol{E}_{\text{total}}(\boldsymbol{r'})d^3\boldsymbol{r'} \tag{2.127}$$

where $G(\boldsymbol{r}, \boldsymbol{r'})$ is the Green's function or the impulse response, which satisfies:

$$|\nabla^2 + k_b^2|G(\boldsymbol{r}, \boldsymbol{r'}) = -\delta(\boldsymbol{r} - \boldsymbol{r'}) \tag{2.128}$$

where $\delta(\boldsymbol{r} - \boldsymbol{r'})$ is the Dirac function.
We now derive the mathematical expression of $\boldsymbol{E}^i(\boldsymbol{r})$. In case of monostatic radars, the incident field in a homogeneous background medium satisfies the Helmholtz equation. For a point source at position $\boldsymbol{r_s}$, the incident field is proportional to the free-space Green's function:

$$\boldsymbol{E}^i(\boldsymbol{r}) = \boldsymbol{E}_0 \cdot G_0(\boldsymbol{r}, \boldsymbol{r_s}) = \boldsymbol{E}_0 \cdot \frac{e^{ik_b|\boldsymbol{r} - \boldsymbol{r_s}|}}{4\pi|\boldsymbol{r} - \boldsymbol{r_s}|} \tag{2.129}$$

where $\boldsymbol{E}_0$ represents the source amplitude. When the far-field approximation ($|\boldsymbol{r} - \boldsymbol{r}_s| \gg \lambda_b$) is valid, this equation describes an outward-propagating spherical wave.
Finally, the scattered field $\boldsymbol{E}^s$ can be written as:

$$\boldsymbol{E}^s(\boldsymbol{r}) = \boldsymbol{E}_{\text{total}}(\boldsymbol{r}) - \boldsymbol{E}^i(\boldsymbol{r}) = \int_V G(\boldsymbol{r}, \boldsymbol{r'})V(\boldsymbol{r'})\boldsymbol{E}_{\text{total}}(\boldsymbol{r'})d^3\boldsymbol{r'} \tag{2.130}$$

Solving Eq.**2.130** for any given specified scattering potential (the formal formulation of the forward scattering problem) leads to a Fredholm integral equation of the second kind, since E_{total} appears both explicitly and inside the integral. When $V(\mathbf{r'})$ is known, this constitutes a well-posed problem in the sense of Hadamard (see Section **2.5.3**), but its numerical solution requires iterative optimization methods, resulting in significant computational cost.

2.5.2 Born Approximation

To clarify the notation used hereafter, we introduce the reflectivity profile $\rho(\boldsymbol{r'})$, defined over the investigation domain Ω, which encodes both the reflectivity value and the spatial geometry of the target (i.e., the spatial points where electromagnetic properties differ from the background). This allows us to rewrite the scattering potential as $V(\boldsymbol{r'}) = k_b^2\rho(\boldsymbol{r'})$, with $\boldsymbol{r'} \in \Omega$. The measurement points $\boldsymbol{r}$ are defined on the measurement domain Γ (the radar acquisition surface).
The *Born approximation* (BA) linearizes the scattering problem by assuming that the scattered field is much weaker than the incident field throughout the scattering volume:

$$|\boldsymbol{E}^s(\boldsymbol{r'})| \ll |\boldsymbol{E}^i(\boldsymbol{r'})|, \quad \boldsymbol{r'} \in \Omega \tag{2.131}$$

Under this assumption, the total field inside the scattering region is approximately equal to the incident field alone $\boldsymbol{E}_{\text{total}}(\boldsymbol{r}') \approx \boldsymbol{E}^i(\boldsymbol{r}')$. From a physical point of view, the BA consists in neglecting the mutual interactions between any two different buried targets or between any two different parts of the same target [63]. Mathematically, under BA assumption, Equation **2.130** becomes:

$$\boldsymbol{E}^s(\boldsymbol{r}) \approx k_b^2 \int_\Omega G(\boldsymbol{r}, \boldsymbol{r}')\rho(\boldsymbol{r}')\boldsymbol{E}^i(\boldsymbol{r}')d^3\boldsymbol{r}', \quad \boldsymbol{r} \in \Gamma \tag{2.132}$$

This equation is now a linear relationship between the reflectivity profile $\rho(\boldsymbol{r}')$ and the scattered field. The forward scattering problem is then:

$$\boldsymbol{E}^s(\boldsymbol{r}) = \mathcal{A}[\rho] \tag{2.133}$$

where $\mathcal{A}$ is the linear forward scattering operator. It is important to distinguish the role of the BA depending on whether it is applied to the forward or the inverse problem. In the forward problem, where $\rho(\boldsymbol{r}')$ is known, the BA provides an accurate approximation of the scattered field only when the scatterers are weak, i.e., when the following quantitative criteria are satisfied:

- weak scattering parameter:

  $$k_b a|\rho(\boldsymbol{r})| \ll 1 \tag{2.134}$$

  where a is the characteristic size of the scatterer (thickness for extended layers, diameter for point scatterers). This condition ensures that the phase perturbation induced by the target is small.

- Small reflectivity:

  $$|\rho(\boldsymbol{r})| \ll 1 \tag{2.135}$$

  This ensures that the fractional change in permittivity is small, minimizing multiple scattering contributions.

A commonly used implementation of these conditions is [76]: $k_b a|\rho(\boldsymbol{r})| < 0.3$, which guarantees that the scattered field amplitude is less than $\sim 30\%$ of the incident field.

In the inverse problem, however, neither $\rho(\boldsymbol{r}')$ nor $\boldsymbol{E}_{\text{total}}$ are known. In this case, the BA is introduced as a linearization strategy: by replacing $\boldsymbol{E}_{\text{total}}$ with $\boldsymbol{E}^i$, the nonlinear and ill-posed inverse problem is reduced to a linear one. This linearization introduces a model error that prevents quantitatively accurate reconstruction of the permittivity profile. However, if the goal is qualitative imaging (i.e., reconstructing the spatial distribution of scatterers without recovering their exact electromagnetic properties), the BA can be applied even when the weak scatterer conditions are not strictly satisfied.

In real applications, like with MARSIS data, we deal with a discrete version of the forward scattering problem. Consider the investigation domain Ω discretized into N voxels at positions $\{\boldsymbol{r}'_n\}_{n=1}^N$, with reflectivity profile values $\rho_n = \rho(\boldsymbol{r}'_n)$. The radar makes M measurements at positions $\{\boldsymbol{r}_m\}_{m=1}^M \in \Gamma$.

Under BA, the forward scattering problem becomes the matrix-vector equation:

$$\boldsymbol{d} = \boldsymbol{A}\boldsymbol{\rho} + \boldsymbol{n} \tag{2.136}$$

where $\boldsymbol{d} \in \mathbb{C}^M$ is the data vector, $\boldsymbol{\rho} \in \mathbb{R}^N$ is the discrete reflectivity profile vector (encoding both scatterer amplitude and geometry at each voxel), $\boldsymbol{A} \in \mathbb{C}^{M\times N}$ is the forward scattering matrix, where M is the number of radar measurements and N is the number of reconstruction voxels, and $\boldsymbol{n} \in \mathbb{C}^M$ represents measurement noise (thermal, clutter, system errors). This matrix equation is the foundation for the inverse problem, where we seek to recover $\boldsymbol{\rho}$ from measurements $\boldsymbol{d}$. an ill-posed inverse problem requiring regularization.

2.5.3 The Inverse Scattering Problem

Under the BA, the inverse scattering problem formulation can be inferred starting from Equation **2.136**: given the observed radar data, the objective is to reconstruct the reflectivity profile of the target:

$$\boldsymbol{\rho} = \boldsymbol{A}^{-1}\boldsymbol{d} \tag{2.137}$$

According to [77], a problem is said to be well posed if it satisfies the following conditions:

- the solution of the problem is unique;
- the solution exist for any data;
- the solution depends with continuity on the data.

The last criterion means that a small (or, mathematically, infinitesimal) error on the data causes a small error in its retrieved solution. In MARSIS context, it means that a small error in the evaluation of the scattered field yields to a small error in the estimation of the complex permittivity value. Unfortunately the inverse scattering problem we are dealing with violates all the previous three conditions:

- failure of uniqueness: when $M < N$, many different solutions exist;
- failure of existence: the presence of noise in the data implies that the data $\boldsymbol{d}$ may not lie in the range of $\boldsymbol{A}$. In addition to that, no exact solution exists and it is necessary to introduce some form of approximation;
- failure of stability: small noise in the data may cause large errors in the solution.

Accordingly, a regularization strategy must be introduced to obtain an approximate and stable solution of Equation **2.137**.
One of the most useful tool in the theory of linear inverse problems is the singular value decomposition (SVD) of a matrix [78]. The standard formulation is the following: *let* $\boldsymbol{A}$ *be a matrix* $M \times N$ *with rank* p*; then there exists a* $p \times p$ *diagonal matrix* Σ*, with positive diagonal elements, and two isometric matrices* $\boldsymbol{U}$ *and* $\boldsymbol{V}$*, respectively* $M \times p$ *and* $N \times p$*, such that:* [74]

$$\boldsymbol{A} = \boldsymbol{U}\boldsymbol{\Sigma}\boldsymbol{V}^* \tag{2.138}$$

where $\boldsymbol{V}^*$ denotes the adjoint of the matrix $\boldsymbol{V}$. If $\boldsymbol{V}$ is a real values matrix: $V^* = V^T$. Moreover, $\boldsymbol{V}$ is the matrix of the singular vectors in the data space, while $\boldsymbol{U}$ is the matrix of the singular vectors in the space of the unknowns, and $\boldsymbol{\Sigma}$ is the diagonal matrix of the singular values.
The forward problem of Equation **2.136** decomposes as:

$$\boldsymbol{d} = \boldsymbol{A}\boldsymbol{\rho} = \sum_{i=1}^{p} \sigma_i(\boldsymbol{v}_i^T\boldsymbol{\rho})\boldsymbol{u}_i \tag{2.139}$$

the vectors $\{\boldsymbol{u_i}\}_{i=1}^{M}$ are the columns of the matrix $\boldsymbol{U}$ and they form an orthonormal basis for the data space. Each element represents the measurement pattern produced by the corresponding reflectivity pattern $\boldsymbol{v}_i$. Similarly, each element of $\{\boldsymbol{v_i}\}_{i=1}^{N}$ represents a spatial reflectivity pattern in the reconstruction volume, ordered by decreasing contribution to the measured signal. Finally, $\{\sigma_i\}_{i=1}^{P}$ are the singular values of $\boldsymbol{A}$. They form the diagonal matrix Σ, such that $\sigma_1 \geq \sigma_2 \geq \sigma_p > 0$. Each σ_i quantifies the sensitivity of the measurements to the reflectivity pattern $\boldsymbol{v}_i$: large σ_i describes a strongly observable feature, while small σ_i a weakly observable one. The rank $p = rank(\boldsymbol{A}) \leq min(M, N)$ is the number of non-zero singular values. Thanks to the SVD approach, the forward scattering problem is decomposed into independent channels: each model pattern $\boldsymbol{v}_i$ maps to the data pattern $\boldsymbol{u}_i$, scaled by σ_i.

By means of the SVD, two distinct approaches can be used to obtain a stable solution of the linearized inverse problem.

The first approach is the adjoint operator $\boldsymbol{A}^*$, which will be described in Section **2.6.2**. The second approach seeks a formal solution to the linear inverse problem. By means of the Moore-Penrose pseudoinverse $\boldsymbol{A}^+$, the least-squares solution is:

$$\boldsymbol{\rho}_{LS} = \boldsymbol{A}^+\boldsymbol{d} = \sum_{i=1}^{p} \frac{1}{\sigma_i}(\boldsymbol{u}_i^*\boldsymbol{d})\boldsymbol{v}_i \tag{2.140}$$

Here the singular values appear at the denominator: components corresponding to small σ_i are amplified, making this solution unstable in the presence of noise. To deal with that, the discrete Picard condition [79] can be applied: for a meaningful solution to exist, the Fourier coefficients of the data must decay faster than the singular values. Since the singular values decay rapidly to zero, only the first $k \ll p$ components contain useful information; beyond k, the pseudoinverse reconstructs noise rather than signal. Regularization stabilizes Equation **2.140** by limiting the contribution of small singular values. In general, the regularized solution can be written as:

$$\boldsymbol{\rho}_{\text{reg}} = \sum_{i=1}^{p} f_i \frac{1}{\sigma_i}(\boldsymbol{u}_i^*\boldsymbol{d})\boldsymbol{v}_i \tag{2.141}$$

where $f_i \in [0, 1]$ are filter factors. The truncated singular value decomposition (TSVD) apply the filter as a hard cutoff. For the chosen value of k:

$$f_i = \begin{cases} 1, & i \leq k \\ 0, & i > k \end{cases} \tag{2.142}$$

$$\boldsymbol{\rho}_{\text{TSVD}} = \sum_{i=1}^{k} f_i \, \frac{1}{\sigma_i} \, (\boldsymbol{u}_i^*\boldsymbol{d}) \, \boldsymbol{v}_i \tag{2.143}$$

By discarding components $i > k$, we lose the ability to reconstruct any reflectivity structure that projects significantly onto $\mathbf{v}_{k+1}, \mathbf{v}_{k+2}, \ldots$. These are typically fine-scale structures. TSVD therefore produces smoother images than the formal pseudoinverse, since the resolution is effectively limited to the spatial scales captured by the first k singular vectors.

It is evident that the outcome of this approach lies in the correct choice of the truncation number k. According to [80], the discrepancy principle states that the optimal value for k is the one for which the data residual matches the noise level:

$$||\mathbf{A}\rho_{\text{TSVD}} - \mathbf{d}|| \approx ||\boldsymbol{n}|| \tag{2.144}$$

The reason behind this choice is clear: for values larger than k noise pattern associated with these fine targets contributes more than the signal itself.

The L-curve method [81] doesn't require knowledge of $||\boldsymbol{n}||$. Figure **2.11** shows the solution norm plotted against the residual norm for $k = 1, 2, ..., p$. For small k, the residual is large (poor data fit) but the solution norm is small (smooth). For large k, the residual is small (good data fit) but the solution norm explodes (noise amplification). The corner of the L shaped curve is the optimal k: the point where further increasing k begins to amplify noise rather than recover signal. The only limitation of this approach lies in the hard cutoff itself. In real data the transition between signal-dominated and noise-dominated is gradual, which means that a component at $i = k + 1$ may still contain useful information to the overall image.

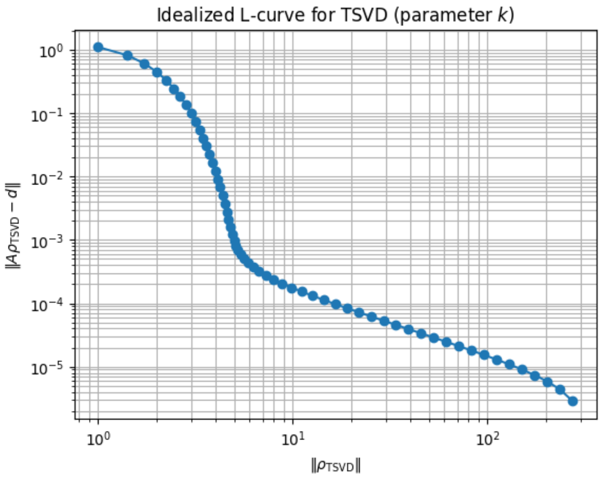

Figure 2.11: L-curve analysis for TSVD

A further strategy to solve the inverse scattering problem defined by Equation **2.137** is the Tikhonov regularization [82], that addresses this problem by introducing a smooth, continuous suppression of small singular values rather than the hard cutoff:

$$\boldsymbol{\rho}_{\text{Tik}} = \arg\,\min\{||\boldsymbol{A\rho} - \boldsymbol{d}||^2 + \alpha||\boldsymbol{L\rho}||^{\mathbf{2}}\} \tag{2.145}$$

where α is the regularization parameter and $\boldsymbol{L}$ the regularization operator, which defines the kind of penalty we apply. In case of $\boldsymbol{L} = 1$ (zeroth-order Tikhonov, minimum norm), the solution becomes:

$$\boldsymbol{\rho}_{\text{Tik}} = \sum_{i=1}^{p} \frac{\sigma_i}{\sigma_i^2 + \alpha}(\boldsymbol{u_i^*d})\boldsymbol{v}_i \tag{2.146}$$

with filter factor $f_i = \frac{\sigma_i^2}{\sigma_i^2+\lambda}$, such that:

- for $\sigma_i \gg \sqrt{\alpha} \quad \rightarrow f_i \approx 1$ the component is inverted normally, as in the pseudoinverse. This holds for well-conditioned components;
- for $\sigma_i \ll \sqrt{\alpha} \quad \rightarrow f_i \approx 0$ the component is completely suppressed. The transition between the two regimes occurs gradually, at $\sigma_i \approx \sqrt{\alpha}$

The regularization parameter α is chosen according the same criteria applied to the TSVD regularization parameter k.

2.6 Migration Algorithms as Approximate Inverses

In the previous Sections we described four practical migration algorithms, that we evaluated for reconstructing 3D volumes of Mars subsurface using MARSIS data. After that, we introduced the mathematical framework of the inverse scattering problem, starting from the definition of the forward scattering problem. This Section aims to be a bridge between the two: in the following pages we will show how each migration algorithm emerges as an approximate solution to the inverse problem under specific assumptions. This Section does not re-describe the algorithms themselves since their implementation and properties have already been covered. Rather, we derive the theoretical foundation of each method, showing which inverse operator it implements and under what conditions the approximation is valid.

Before that, we must first establish the central role of the Green's function - the mathematical object that bridges abstract operator theory and concrete formulas.

2.6.1 The Green's Function as the Forward Operator Kernel

We recall that, under the BA, the Lippmann-Schwinger equation (Equation **2.130**) is the integral representation of the forward scattering problem. The Green's function $G(\boldsymbol{r}, \boldsymbol{r}')$ is also called the propagator, since it describes how a scatterer located at $\boldsymbol{r}'$ contributes to the field at the observation point $\boldsymbol{r}$. The Green's function is the fundamental solution of the wave equation, encoding all propagation physics (geometric spreading, phase accumulation, and polarization). Hence, the forward operator that we introduced in Equation **2.136** can now be written explicitly:

$$\mathbf{A}\rho = \int_V G(\mathbf{r}, \mathbf{r}')E^i(\mathbf{r}')k_b^2\rho(\mathbf{r}')d^3\mathbf{r}' \tag{2.147}$$

$$\mathcal{A}(\mathbf{r}, \mathbf{r}') = k_b^2 G(\mathbf{r}, \mathbf{r}')E^i(\mathbf{r}') \tag{2.148}$$

The matrix $\mathbf{A}$ is a concrete integral operator, with its kernel $\mathcal{A}$ determined by Green's function G, the incident field E^i and the wavenumber factor k_b^2.

In case of a radar operating in a homogeneous background medium (relative permittivity ϵ_r), the Green's function of a point scatterer is defined as in Equation **2.129**. In a monostatic configuration, the forward operator kernel therefore involves the product of the square Green's function:

$$\mathcal{A}(\mathbf{r}_s, \mathbf{r}') = k_b^2 G(\mathbf{r}_s, \mathbf{r}')E^i(\mathbf{r}') = \frac{E_0 k_b^2}{(4\pi)^2 R^2} e^{j2k_b R} \tag{2.149}$$

with $R = |\mathbf{r}_s - \mathbf{r}'|$. By using $\tau = 2R/c$ and $k_b = \omega/c$, we can rewrite the kernel as the explicit frequency-domain forward operator for monostatic radar:

$$\mathcal{A}(\mathbf{r}_s, \mathbf{r}', \omega) = \frac{E_0 k_b^2}{(4\pi)^2 R^2} e^{j\omega\tau(\mathbf{r}, \mathbf{r}')} \tag{2.150}$$

Thus, the forward problem for a single frequency is:

$$E^s(\mathbf{r}_s, \omega) = \int_V \mathbf{A}(\mathbf{r_s}, \mathbf{r}', \omega)\rho(\mathbf{r}')d^3\mathbf{r}' \tag{2.151}$$

2.6.2 The Adjoint Operator and TD Backprojection

In Section **2.5.3** we introduced the pseudoinverse matrix as a way to solve the linearized inverse scattering problem. An alternative could be using the adjoint operator $\mathbf{A}^*$, which is defined as the complex conjugate of the forward operator $\mathbf{A}$. Using the SVD approach, the adjoint can be written as:

$$\mathbf{A}^* = (\mathbf{U}\mathbf{\Sigma}\mathbf{V}^*)^* = \mathbf{V}\mathbf{\Sigma}\mathbf{U}^* \tag{2.152}$$

The adjoint maps from data space back to model space:

$$\rho_{\text{adj}} = \sum_{i=1}^{p} \sigma_i(\boldsymbol{u}_{\boldsymbol{i}}^{*}\boldsymbol{d})\boldsymbol{v}_i \tag{2.153}$$

Note that in this expression the singular values σ_i appear at the numerator: components associated with small singular values are naturally suppressed, making the adjoint operator inherently stable. However, Equation **2.153** does not provide a formal solution to the inverse problem: even in the absence of noise, the true reflectivity profile cannot be exactly recovered. From a practical standpoint, this poses an intrinsic resolution limit independent of the noise level [83].

The adjoint, the TSVD, and Tikhonov regularization thus represent two fundamentally different strategies: the adjoint avoids instability by construction (at the cost of not achieving a formal inversion), whereas TSVD and Tikhnov regularize the formal inverse by truncating the expansion in different ways. A systematic comparison between the adjoint-based reconstruction and the TSVD approach, carried out in terms of achievable resolution limits via point spread function analysis, is provided in [83]. As shown in that work, the two approaches achieve comparable resolution when the investigation domain is large, while TSVD outperforms the adjoint when the domain can be constrained.

Mathematically, the adjoint operator $\mathbf{A}^*$ is defined as the conjugate transpose of $\mathbf{A}$. In a matrix approximation, this corresponds to taking the Hermitian transpose $\mathbf{A}^* = (\mathbf{A}^T)^\dagger$. The kernel of the adjoint operator at angular frequency ω is obtained by conjugating the kernel of Eq. **2.150**:

$$\mathcal{A}^*(\mathbf{r}_s, \mathbf{r}', \omega) = \frac{E_0 k_b^2}{(4\pi)^2 R^2} e^{+j\omega\tau(\mathbf{r},\mathbf{r}')} \tag{2.154}$$

where $\mathcal{A}^*$ denotes the kernel of $\mathbf{A}^*$, distinct from the operator itself. The change $e^{-j\omega\tau} \to e^{+j\omega\tau}$ represents a time-reversed wave propagating back toward the source.

Applying this spatial adjoint to the data at frequency ω gives:

$$\tilde{\rho}(\mathbf{r}', \omega) = \int_S \frac{1}{R^2(\mathbf{r}_s, \mathbf{r}')} E^s(\mathbf{r}_s, \omega) e^{-j\omega\tau(\mathbf{r},\mathbf{r}')dS_s} \tag{2.155}$$

where the integral is over the acquisition surface S (all radar positions). The physical reflectivity ρ is then given by:

$$\rho(\mathbf{r}') = \int_{-\infty}^{\infty} \left[\int_S \frac{1}{R^2(\mathbf{r}_s, \mathbf{r}')} E^s(\mathbf{r}_s, \omega) e^{-j\omega\tau(\mathbf{r},\mathbf{r}')dS_s} \right] d\omega \tag{2.156}$$

The full adjoint operator includes both spatial backpropagation and integration over frequency, corresponding to an inverse Fourier transform:

$$\int E^s(\mathbf{r}_s, \omega) e^{-i\omega\tau} d\omega = s(\tau, \mathbf{r}_s) \tag{2.157}$$

Now we can derive the final solution as:

$$\rho(\mathbf{r}') = \int_S \frac{1}{R^2} s(\tau(\mathbf{r}_s, \mathbf{r}'), \mathbf{r}_s)\, dS_s \tag{2.158}$$

Which is the time-domain backprojection formula of Equation **2.109**.

2.6.3 Frequency-Domain Migration Algorithms

The time-domain backprojection algorithm implements the adjoint operator A* under minimal assumptions (BA and homogeneous background). However, for certain simplified geometries, the forward operator A takes a special form that allows a closed-form inversion rather than adjoint application. Stolt migration (Section **2.4.3**) corresponds to another closed-form inversion of the linearized forward operator because, under constant-velocity and flat-geometry assumptions, the operator becomes diagonal in the Fourier domain.

The exploding reflector model combined with constant velocity c and flat surface $z = 0$ produces a forward operator A with special structure. In the frequency-wavenumber domain (k_x, ω), the relationship between data $\tilde{d}(k_x, \omega)$ and reflectivity spectrum $\tilde{\rho}(k_x, k_z)$ becomes:

$$\tilde{d}(k_x, \omega) = H(k_x, \omega) \cdot \tilde{\rho}(k_x, k_z(\omega, k_x)) \tag{2.159}$$

where H is the propagation factor and k_z is given by the dispersion relation of Equation **2.118**. This operator is diagonal: each component (k_x, ω) evolves independently, with no coupling between different wavenumber-frequency pairs. Thus, in this scenario it is possible to compute the inverse matrix as:

$$\mathbf{A}^{-1} = \text{diag}\left\{ \frac{1}{H(k_x, \omega)} \right\} \tag{2.160}$$

Geometrically, Stolt mapping re-indexes data from the measurement coordinates (k_x, k_z) to image coordinates (x,z) dividing by the singular values of the propagation operator. We stress again that this approach is valid only if all the assumptions are met: if one of them fails, phase errors accumulate, producing artifacts that can be larger than the true signal.

Omega-k algorithm described in Section **2.4.3** extends Stolt's approach to handle bistatic configurations. Hence, as we have seen before, omega-k implement the closed-form inversion A^{-1} if all the hypothesis are met.

Phase-shift migration (Section **2.4.3**) differs fundamentally from Stolt and omega-k: rather than attempting global spectral inversion, it implements the adjoint operator A^* through iterative backward wave propagation. This connects it conceptually to time-domain backprojection while retaining frequency-domain computational efficiency.

The forward problem (upward propagation under exploding reflector model assumption) satisfies:

$$\frac{\partial \tilde{P}}{\partial z} = -jk_z(z)\tilde{P} \tag{2.161}$$

The solution for upward propagation through depth Δz is:

$$\tilde{P}(z - \Delta z) = \tilde{P}(z) \cdot e^{-jk_z(z)\Delta z} \tag{2.162}$$

The adjoint reverses the propagation direction (backward/downward) by changing the sign:

$$\tilde{P}(z - \Delta z) = \tilde{P}(z) \cdot e^{+jk_z(z)\Delta z} \tag{2.163}$$

This is exactly the phase-shift operator (Equation **2.119**). By applying it iteratively from the surface ($z = 0$) downward, phase-shift migration implements:

$$\rho = \mathbf{A}^* \mathbf{d} \tag{2.164}$$

Although similar to time-domain backprojection, phase-shift migration still requires flat surface and smooth velocity assumptions.

Table **2.6** provides a schematic summary of all the algorithms we described so far.

Table 2.6: Comparison of imaging algorithms

Algorithm	**Operator**	**Assumptions**	**Accuracy**
Time-domain BPA	$\mathcal{A}^*$	BA only	Positioning: excellent Amplitude: poor
Stolt	$\mathcal{A}^{-1}$	BA, flat surface, constant velocity c, zero-offset	Positioning: excellent Amplitude: excellent
Omega-k	$\mathcal{A}^{-1}$	BA, smooth velocity $c(x, z)$, known velocity model	Positioning: excellent Amplitude: good
Phase-shift	$\mathcal{A}^*$	BA, layered velocity $c(z)$, one-way wave assumption	Positioning: excellent Amplitude: good

2.6.4 Point Spread Function: Visualizing the Difference

The abstract distinction between inverse operators (A^{-1}, A^+) and the adjoint (A^*) becomes concrete when we examine how each reconstructs a single point scatterer. This analysis reveals why adjoint-based methods like backprojection produce "blurred" images while true inversion (when achievable) yields sharp reconstruction.

Consider a point scatterer embedded in a homogeneous medium. The true reflectivity is described by the Dirac function as:

$$\rho_{\text{true}}(x, z) = \delta(x - x_0, z - z_0) \tag{2.165}$$

The associated forward scattering problem is:

$$\mathbf{d} = \mathbf{A}[\delta(x - x_0, z - z_0)] \tag{2.166}$$

As we saw in the previous sections, the data vector is composed by echoes collected at each radar position, received after a travel time τ corresponding to its distance from the scatterer location. Thus, **d** consists of hyperbolic echoes in the range-azimuth domain.

If the forward operator **A** is square (M = N) and full rank, it is invertible (for instance, when the Stolt's assumptions are valid). Thus, it is possible to recover the target reflectivity as:

$$\rho_{rec} = \mathbf{A^{-1}d} = \mathbf{A}^{-1}\mathbf{A}[\delta] = \delta(x - x_0, z - z_0) = \text{PSF}_{\mathbf{A}^{-1}} \tag{2.167}$$

This is the response of an ideal imaging system: the inverse problem focused the data at the target location (x_0, z_0) without spreading, sidelobes or geometric artifacts.

When the forward problem is underdetermined (M < N), the matrix A is rectangular. In this case the true inverse does not exist, but the pseudoinverse A^+ provides the minimum-norm least-squares solution.

Starting from the SVD decomposition of the forward problem applied to a point scatterer $\delta(x - x_0, z - z_0)$:

$$\mathbf{A}[\delta] = \sum_{m=1}^{p} \sigma_m(\mathbf{v}_m^T\delta)\mathbf{u}_m = \sum_{m=1}^{p} \sigma_m v_m(x - x_0, z - z_0)\mathbf{u}_m \tag{2.168}$$

where $v_m(x - x_0, z - z_0) = \mathbf{v}_m^T \delta$ is the m-th singular vector evaluated at the scatterer position. Substituting into the TSVD operator:

$$\mathrm{PSF}_{\mathrm{TSVD}} = \left[\sum_{i=1}^{k} \frac{1}{\sigma_i} \mathbf{v}_i \mathbf{u}_i^*\right] \left[\sum_{m=1}^{p} \sigma_m v_m(x - x_0, z - z_0) \mathbf{u}_m\right] \tag{2.169}$$

$$= \sum_{i=1}^{k} \sum_{m=1}^{p} \frac{\sigma_m}{\sigma_i} v_m(x - x_0, z - z_0)\, \mathbf{v}_i \,(\mathbf{u}_i^* \mathbf{u}_m) \tag{2.170}$$

Applying the orthonormality of the left singular vectors, $\mathbf{u}_i^* \mathbf{u}_m = \delta_{im}$, the double sum collapses:

$$\mathrm{PSF}_{\mathrm{TSVD}} = \sum_{i=1}^{k} \frac{\sigma_i}{\sigma_i} v_i(x - x_0, z - z_0)\, \mathbf{v}_i = \sum_{i=1}^{k} v_i(x - x_0, z - z_0)\, \mathbf{v}_i = \sum_{i=1}^{k} \mathbf{v}_i \mathbf{v}_i^T\, \delta \tag{2.171}$$

which corresponds to the projection of the delta function onto the first k right singular vectors, confirming Equation **2.143**.

If we apply the Tikhonov regularization (Equation **2.146**):

$$\mathrm{PSF}_{\mathrm{Tik}} = \sum_{i=1}^{p} \frac{\sigma_i^2}{\sigma_i^2 + \alpha} \mathbf{v}_i \mathbf{v}_i^T \delta \tag{2.172}$$

The resolution (and thus, the presence of sidelobes) depends on the regularization parameters: the smaller is k (larger α), the wider is the PSF. On the other hand, large k (small α) corresponds to a narrower PSF. However, in this case the solution is less stable (much more sensitive to noisy patterns). The choice of the regularization parameter requires a trade-off between resolution and stability.

If we solve the inverse scattering problem by means of the adjoin operator (as in time-domain backprojection), without explicitly attempting inversion:

$$\rho_{\mathrm{BP}} = \mathbf{A}^* \mathbf{d} = \mathbf{A}^* \mathbf{A}[\delta] \tag{2.173}$$

The product A*A is the Gram matrix:

$$\mathbf{A}^* \mathbf{A} = \sum_{i=1}^{r} \sigma_i^2 \mathbf{v}_i \mathbf{v}_i^T \tag{2.174}$$

And the solution becomes:

$$\mathrm{PSF}_{\mathbf{A}^*} = [\mathbf{A}^* \mathbf{A}]\delta = \sum_{i=1}^{r} \sigma_i^2 (\mathbf{v}_i^T \delta) \mathbf{v}_i \tag{2.175}$$

As we have already stated, the adjoint operator acts as a low-pass filter. Thus, the PSF is the autocorrelation of the antenna pattern convolved with the signal spectrum, representing the natural resolution limit imposed by wavelength, bandwidth, and aperture size.

2.7 Tomographic Imaging and Multi-Angle Reconstruction

In the previous sections we established the mathematical and algorithmic foundations of radar imaging. The purpose of this section is to place these methods in the broader context of 3D tomographic reconstruction, where measurements acquired from multiple, geometrically diverse viewing angles are combined to reconstruct a three-dimensional volume. This perspective reveals both the power and fundamental limitations of radar sounding for subsurface imaging.

Tomography algorithms are adopted in different fields: medical imaging (computed tomography, CT), seismology (earthquake wave inversion), electron microscopy (cryo-EM), and radio astronomy (interferometry). The mathematical foundation of the algorithm is always the same in any of these context. In our case, we will explore how the limited viewing geometry affects the reconstruction quality of the final image. To this purpose, we will introduce the classical Projection-Slice Theorem of geometrical tomography to account for wave propagation effects.

2.7.1 The Tomographic Inverse Problem

Tomographic reconstruction is the inverse problem of recovering a two- or three-dimensional scalar field $f(\mathbf{r})$ from a set of measurements acquired along different paths or from different viewing angles ([84, 85]).

We start the derivation of the tomography algorithm by recalling the BA of Equation **2.132**. Consider a monostatic radar at altitude h above the surface. The incident field can be written as:

$$E_i(\mathbf{r}') = E_0 e^{j\mathbf{k}_{\text{inc}}\mathbf{r}'} \tag{2.176}$$

for a nadir-looking radar, the incident wavevector is $\mathbf{k}_{\text{inc}} = k_0\hat{z}$. We also recall that the Green's function for a point scatterer is defined in Equation **2.129**. In the far field approximation, the distance between the subsurface point and the receiver $|\mathbf{r} - \mathbf{r}'|$ is:

$$|\mathbf{r} - \mathbf{r}'| \approx R - \hat{\mathbf{k}}_{\text{scat}} \cdot \mathbf{r}' \tag{2.177}$$

where $R = h$ and $\hat{\mathbf{k}}_{\text{scat}}$ is the unit vector defining the scattering direction $\mathbf{k}_{\text{scat}} = k_0\hat{\mathbf{k}}_{\text{scat}}$. Substituting into the Green's function:

$$G(\mathbf{r}, \mathbf{r}') \approx \frac{e^{jk_0R}}{4\pi R} e^{-j\mathbf{k}_{\text{scat}}\mathbf{r}'} \tag{2.178}$$

Thus, the BA becomes:

$$E_s(\mathbf{r}) \approx \frac{k_b^2 E_0 e^{jk_0R}}{4\pi R} \int_V \rho(\mathbf{r}') e^{j(\mathbf{k}_{\text{inc}} - \mathbf{k}_{\text{scat}})\mathbf{r}'} d^3\mathbf{r}' \tag{2.179}$$

By defining the scattering wavevector as:

$$\mathbf{K} = \mathbf{k}_{\text{inc}} - \mathbf{k}_{\text{scat}} \tag{2.180}$$

Then:

$$E_s(\mathbf{r}) \propto \int_V \rho(\mathbf{r}') e^{j\mathbf{K}\mathbf{r}'} d^3\mathbf{r}' = \tilde{\rho}(\mathbf{K}) \tag{2.181}$$

Thus, the scattered field measured by the radar is the Fourier transform of the subsurface reflectivity (up to a constant), evaluated at the specific frequency $\mathbf{K}$. This statement is known as the Fourier Diffraction Theorem [86], which can be formally stated as: *Each radar measurement at frequency ω and scattering geometry ($\mathbf{k}_{inc}, \mathbf{k}_{scat}$) provides one sample of the 3D Fourier transform of the subsurface at spatial frequency $\mathbf{K} = \mathbf{k}_{inc} - \mathbf{k}_{scat}$.* This geometrical relation is presented in Figure **2.12**. The subsurface reflectivity $\rho(\mathbf{r}')$ could be fully reconstructed by measuring $\tilde{\rho}(\mathbf{K})$ at all spatial frequencies, and by applying the inverse Fourier transform:

$$\rho(\mathbf{r}') = \frac{1}{(2\pi)^3} \int \tilde{\rho}(\mathbf{K}) e^{-j\mathbf{K}\mathbf{r}'} d^3\mathbf{K} \tag{2.182}$$

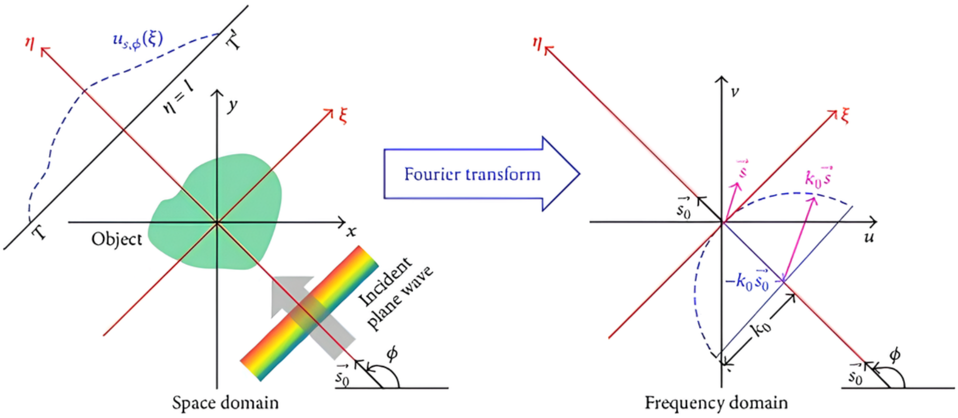

Figure 2.12: Geometrical representation of the Fourier Diffraction Theorem. From [87].

In practical applications, the radar can access only a limited range of **K** values, determined by its operating frequency and viewing geometry. The regions of k-space (the space of spatial frequencies **K**) that remain unsampled correspond to subsurface structures that cannot be reconstructed. For a nadir-looking radar the incident and scattered wavevectors are antiparallel, such that: $\mathbf{K} = 2k_0\hat{z}$. By changing the geometry of the observation or, equivalently, as the viewing angle varies, **K** traces out a spherical surface in k-space with radius $2k_0$, known as the *Ewald Sphere* [88]. Figure **2.13** shows the Ewald Sphere in the k-space. Every measurement made by the radar at frequency ω samples exactly one point on this sphere, dependent on the viewing direction. As the radar changes its look angle, it samples different points on the sphere's surface. Limited angular diversity, such as parallel orbits, leaves large gaps in k-space, resulting in anisotropic resolution. Since every radar operates over a bandwidth $B = f_{\max} - f_{\min}$, the wavenumber varies as:

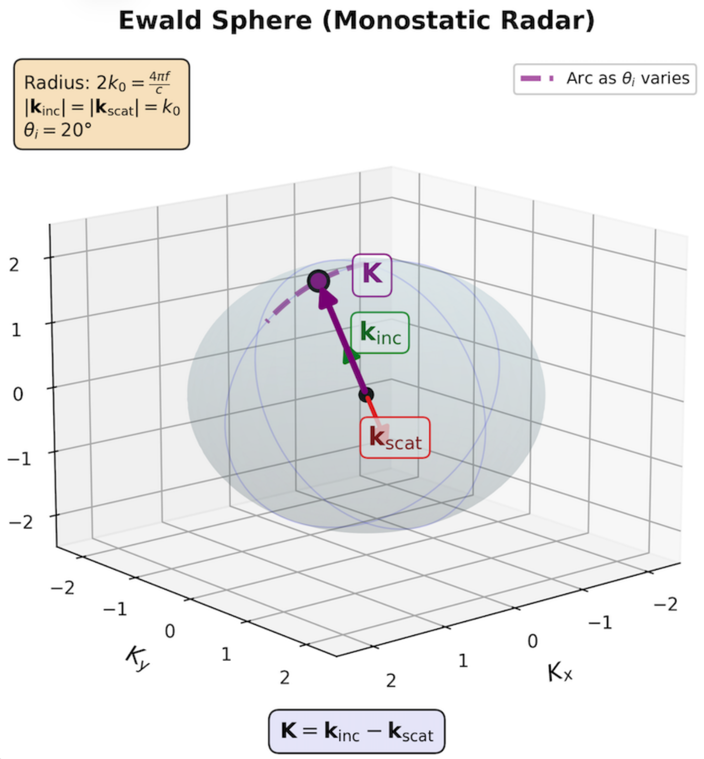

Figure 2.13: The Ewald Sphere in the k-space.

$$k_0 \in \left[\frac{2\pi f_{\text{min}}}{c}, \frac{2\pi f_{\text{max}}}{c}\right] \tag{2.183}$$

This produces a family of concentric Ewald spheres, one for each frequency in the band. The accessible region in k-space is then a spherical shell with inner radius $2k_{min}$ and outer radius $2k_{max}$. The thickness of this shell is:

$$\Delta K = 2(k_{\text{max}} - k_{\text{min}}) = \frac{4\pi B}{c} \tag{2.184}$$

The Fourier Diffraction Theorem is a generalization of the Projection-Slice Theorem [89], which holds when wave effects are neglected and only geometrical ray path are considered. The theorem states that *the 1D Fourier transform of a projection (line integral) through an object at θ equals a radial slice through the object's Fourier transform at angle θ in k-space.*. Mathematically, for a 2D object $f(x, y)$:

$$P_\theta(r) = \int_{-\infty}^{\infty}\int_{-\infty}^{\infty} f(x, y)\delta(r - x \cos\theta - y \sin\theta)dxdy \tag{2.185}$$

The Fourier transform satisfies:

$$\mathcal{F}_r\{P_\theta(r)\}(k) = \mathcal{F}_{xy}\{f(x, y)\}(k \cos\theta, k \sin\theta) \tag{2.186}$$

This is the geometrical optics limit of the Fourier Diffraction Theorem. When subsurface features are much larger than the wavelength, the two approaches give nearly identical results.

2.7.2 Backprojection in the tomographic context

The backprojection algorithm described in Section **2.4.2** is an inversion approach used for tomographic reconstruction. Hence, the quality of the 3D reconstruction depends critically on the angular diversity of the input data: regardless of the inversion strategy adopted, the achievable resolution is ultimately limited by the angular coverage of the acquisitions. In general we can define three distinct regimes according to the degree of angular diversity available:

- in a single orbit configuration, all radar positions lie along a single line with approximately constant altitude h. For a subsurface target at depth z and along-track offset x, the incidence angle from different radar positions varies as:

  $$\theta_i(x) \approx \arctan\left(\frac{x}{h + z}\right) \tag{2.187}$$

  As the radar moves along its orbit over a distance L, the incidence angle varies as:

  $$\Delta\theta \approx \frac{L}{h + z} \tag{2.188}$$

  For typical orbital MARSIS parameters ($h = 300km$, $z = 1$ km and $L = 10$ km), this results in $\Delta\theta < 2°$. With such a narrow angular aperture, the scene is observed from essentially the same viewing angle at all times. As a consequence, any object in the scene responds as if it were a planar interface, and it is not possible to distinguish its actual 3D geometry. The hyperbolic signatures visible in a single-orbit image are therefore a manifestation of the finite cross-section of the target in the range along-track plane, not of its three-dimensional structure.

- For multiple orbits with different ground tracks, separated by distances $\Delta y \sim 30 - 100$ km, the viewing geometry varies between each orbit. Suppose we have M orbits, each separated by cross-track distance Δy. For a subsurface target at position (x_0, y_0, z_0), the cross-track incidence angle from orbit m is:

  $$\theta_y^{(m)} \approx \arctan\left(\frac{y_m - y_0}{h + z_0}\right) \tag{2.189}$$

  where y_m is the cross-track position of the m-th orbit. If the total cross-track range is $\Delta Y = M \cdot \Delta y$, the angular diversity is:

  $$\Delta\theta_y \approx \frac{\Delta Y}{h + z_0} \tag{2.190}$$

  For $m = 3$, $\Delta y = 50$ km and $z_0 = 1$ km: $\Delta\theta_y \approx 23°$. This configuration allows to infer partially the 3D subsurface structure, leaving gaps in the regions without angular coverage. In the k-space, each orbit contributes measurements along a different arc on the Ewald sphere. In this case, multi-orbit backprojection achieves quasi-tomographic imaging: partial 3D reconstruction with highly anisotropic resolution.

- To achieve near-isotropic 3D resolution comparable to medical CT or seismic tomography, large angular diversity is required: $\Delta\theta \sim 30° - 90°$. This can in principle be achieved in a monostatic configuration, provided that multiple orbits with sufficiently different viewing geometries are available. In this case, all measurements sample points on the surface of the Ewald sphere, since $\mathbf{K} = 2k_0\hat{z}$ for nadir-looking geometry. In a bistatic configuration, where transmitter and receiver are on separate platforms, the transmit and receive wavevectors have different orientations, giving access also to the interior of the Ewald sphere. This provides a fundamentally richer k-space coverage, potentially improving the tomographic reconstruction quality beyond what is achievable with a monostatic system.

Chapter 3

MARSIS subsurface sounding of Mars

This Chapter presents the application of the signal processing and imaging techniques developed in Chapter **2** to MARSIS radar sounding data, with the goal of investigating the subsurface structure of the Martian south polar region.

The Chapter is organized as follows. We begin with a brief introduction to the geological and geophysical context of Mars, with particular attention to the history of water and the evidence for past and present liquid water on the planet (Section **3.1**). In Section **3.2**, we describe the MARSIS instrument, its orbital configuration, operating modes, and data acquisition geometry. The study area, the Ultimi Scopuli region of the South Polar Layer Deposits, is introduced in Section **3.3**.

The second part of the chapter is dedicated to data processing and results. In Section **3.4** we describe the MARSIS data processing pipeline, from raw data to range compression and surface simulation and we discuss the main challenges encountered, including high noise levels, residual phase delays, and ionospheric distortions in Section **3.5**. A noise reduction algorithm is presented in Section **3.6**. Finally, we present the 2D and 3D focusing results obtained by applying the frequency-domain backprojection algorithm to a selection of MARSIS orbits in the vicinity of the bright basal reflector in Sections **3.7** and **3.8**.

3.1 Introduction to Mars

3.1.1 Geological and geophysical context

Mars, the fourth planet from the Sun, represents one of the most extensively studied terrestrial bodies in our Solar System and the primary target for comparative planetology with Earth. The fundamental physical and orbital parameters of Mars are summarized in Table **3.1**. Currently, Mars is an harsh and hostile environment, that could not meet the necessary conditions to host life. Its climate is characterized by extremely low mean temperatures (~ 210 K) with significant diurnal variations (up to 100 K) and seasonal cycles driven by the planet's orbital eccentricity and axial obliquity. The thin atmosphere, with a mean surface pressure of ~ 600 Pa, is composed primary of CO_2 and other minor constituents such as N_2 (2.7%), Ar (1.6%), and traces of O_2, CO, and H_2O [90, 91]. This low atmospheric pressure places the triple point of water near or above the surface pressure at most elevations, making liquid water thermodynamically unstable under current surface conditions.

The atmosphere supports a dynamic seasonal cycle involving the formation and sublimation of polar ice caps. The permanent polar caps, composed primarily of water ice and covered by a layer

Table 3.1: Comparison of physical, orbital, and geophysical properties of Mars and Earth.

Physical Property	**Mars**	**Earth**
Bulk Parameters		
Mean radius	3,389.5 km	6,371.0 km
Mass	6.417×10^{23} kg	5.972×10^{24} kg
Mean density	3,933 kg/m^3	5,514 kg/m^3
Surface gravity	3.71 m/s^2	9.81 m/s^2
Escape velocity	5.03 km/s	11.19 km/s
Orbital Characteristics		
Semi-major axis	1.524 AU	1.000 AU
Orbital period	687 Earth days	365.25 days
Orbital eccentricity	0.0934	0.0167
Rotation period	24 h 37 m 22 s	23 h 56 m 04 s
Axial tilt	25.19°	23.44°
Thermal Properties		
Mean surface temperature	210 K (-63°C)	288 K ($+15$°C)
Temperature range	130–300 K	184–330 K
Solar constant	589 W/m^2	1361 W/m^2
Atmospheric Properties		
Surface pressure	$\sim$600 Pa	101,325 Pa
Atmospheric mass	2.5×10^{16} kg	5.1×10^{18} kg
Main constituent	95.3% CO_2	78% N_2
Scale height	$\sim$11 km	$\sim$8.5 km
Internal Structure		
Core radius	$\sim$1830 km	$\sim$3485 km
Mantle thickness	$\sim$1560 km	$\sim$2890 km
Mean crustal thickness	30–70 km	$\sim$40 km
Geophysical Properties		
Active magnetic field	No	Yes
Plate tectonics	No	Yes
Volcanic activity	Extinct / dormant	Active
Largest volcano height	21 km (Olympus Mons)	9 km (Mauna Kea)

of CO_2 ice at the South Pole, constitute the planet's largest known water reservoirs. During winter in each hemisphere, seasonal CO_2 ice deposits accumulate, containing up to $25-30\%$ of the total atmospheric mass, which then sublimes during spring and summer [92]. This massive seasonal exchange of CO_2 between the atmosphere and polar caps drives global atmospheric circulation patterns and pressure variations of $\sim 30\%$.

The Martian surface shows an evident hemispheric dichotomy extensively studied in planetology. The southern hemisphere is composed by heavily cratered highlands, elevated approximately 5 km above the northern lowlands and covering roughly 60% of planet's surface. The analysis of these craters dated them back at the Noachian period (> 3.7 Gy), at the time of early bombardment era of the Solar System. In contrast, the northern lowlands present a relatively smooth, younger surface with fewer impact craters, lower elevation, and potentially modified by ancient ocean or sedimentary processes [93, 94]. The origin of this dichotomy remains one of the most outstanding questions in Martian geology, with hypotheses ranging from a giant impact during early planetary accretion [95] to endogenic processes involving mantle convection [96].

Mars surface possesses several prominent volcaninc and tectonic features. The Tharsis volcanic region, a massive topographic rise centered near the equator, hosts the solar system's largest volcanoes, including Olympus Mons (21 km height, 600 km diameter at base) and the Tharsis Montes [94]. The province's formation, spanning much of Martian history from the Noachian to recent geological times, involved the production of volcanic materials and caused significant crustal deformation [97]. In proximity of Tharsis, the Valles Marineris canyon system extends over 4000 km in length, reaching depths of 7 km and widths up to 200 km, representing a complex tectonic and erosional feature that could be related to Tharsis uplift and subsequent collapse [98].

As for the internal structure of the planet, the InSight mission $(2018-2022)$ provided the first direct seismic measurements of the Martian interior, revealing a partially molten core with a radius of approximately $1,830 \pm 40$ km [99]. The discovery of a liquid outer core gives us hints about the planet thermal evolution and magnetic field history. Topography data collected by Mars Global Surveyor and Mars Reconnaissance Orbiter show that the Martian crust exhibits significant lateral variations in thickness, with estimates ranging from as thin as 10 km beneath the Hellas impact basin to over 100 km beneath the Tharsis volcanic province [100, 101].

Unlike Earth, Mars lacks active plate tectonics, representing instead a "stagnant lid" regime where the lithosphere remains as a single, unbroken shell [102]. This tectonic behavior has profoundly influenced the planet's geological evolution, concentrating volcanic activity in long-lived centers rather than distributing it along plate boundaries. The planet's ancient magnetic field, recorded in strongly magnetized crustal rocks primarily in the southern highlands, indicates the existence of an active dynamo during the first $500-800$ million years of Martian history [103, 104]. The subsequent cessation of this dynamo, likely due to core cooling and solidification, left Mars without the protective magnetosphere that shields Earth from solar wind erosion, with critical consequences for atmospheric retention and surface habitability [105].

Mars possesses two small, irregularly shaped natural satellites. Phobos, the larger and innermost moon, orbits at only 9376 km from the Martian surface (2.76 Mars radii) with a mean radius of 11.3 km, while Deimos orbits at 23436 km with a mean radius of 6.2 km [106]. Both moons are tidally locked to Mars and exhibit very low densities ($\sim 1,900$ kg/m^3 for Phobos, $\sim 1,750$ kg/m^3 for Deimos), suggesting highly porous interiors or composition dominated by volatile-rich materials. Their origin is still unclear and debated: hypothesis span from captured asteroids from the outer main belt to debris from a giant impact early in Mars' history [107]. Phobos is particularly notable for its remarkably low orbit, which causes it to gradually spiral inward due to tidal forces; it is predicted to either collide with Mars or be torn apart by tidal stresses within approximately 50 million years [108].

3.1.2 History of water on Mars

Geologically, observations of the density of craters in a certain region can be used as a means of estimating that region's absolute age [109]. This relation comes from the observation of impact basins on the Moon, whom age can be calculated with high precision.
However, this technique has a great degree of uncertainty when it is applied to Mars, or more in general to other celestial bodies. The reason is straightforward: the number of craters that we count today could not correspond to the overall number of impacts through time, given the superposition of new surface features (such as lava flows, or successive impacts) that could have erased their presence. Nevertheless, researchers managed to estimate that the age limits of the major geological eras on Mars are as follows:

- Noachian era, $\sim$ 4.1 billions to $\sim$ 3.7 billion years ago;
- Hesperian era, $\sim$ 3.7 billions to $\sim$ 3 billions years ago;
- Amazonian era, $\sim$ 3 billions of years ago to the present day.

The Noachian period is characterized by heavy meteorite bombardment, which ended around 3.8 billion years ago. Most surfaces that date from this era are heavily dissected by networks of valleys a few kilometers wide but up to a few thousand kilometers long. Relatively high drainage densities suggest surface runoff that would require either rainfall or melting snow, which in turn implies significantly warmer and wetter conditions than those that prevail today. As pointed out in [98], there are several elements suggesting that Mars was warmer in the Noachian that it is today: the finding of evaporites at the Mars Exploration Rover landing site on Meridiani Planum and elsewhere, evidences for fluctuation of the water table at Meridiani, detection of clay minerals from orbit in Noachian terrains, the presence of hydrothermally altered rocks in the Columbia Hills at the Mars Exploration Rover landing site in Gusev crater, and surface erosion rates that were 4 to 5 orders of magnitude higher than they were subsequently. These observations, in combination, suggest an Earth-like, active hydrological cycle with large lakes or oceans that acted as evaporative sources, sinks, and base levels for erosion. However, Noachian terrains have been mainly shaped by impacts, volcanism, and deformation rather than by fluvial processes. Such global catastrophes could have damaged or canceled any emerging life (most of the liquid water could have been boiled away). However, these extreme events could have been separated by millions of years in which the condition for the sustain and growth of life could have been quietly met.
It is still unclear why Mars was warmer at that age. Surely, evidences of an ancient global magnetic field could have allowed the planet to form a thicker atmosphere, screening the surface from solar UV radiation. In any case, at the end of the Noachian, the rate of formation of valley networks declined rapidly, although not to zero, erosion rates fell precipitously, and clay minerals appear to have stopped forming.
The most characteristic fluvial feature of the post-Noachian era is the outflow channel, formed by large floods, rather than the valley networks that characterize the preceding era. Nevertheless, young valley networks are found in places, such as on young volcanoes, indicating that conditions

necessary for slow erosion by running water were occasionally and locally met. The largest flood channels can be found around the Chryse basin, which are mostly Hesperian in age (some 3.0 to 3.7 billions years ago), but some of them could be as young as a few tens of millions of years, meaning that they could also be formed today [110]. These channels seem to have formed by eruptions of groundwater from below a thick cryosphere or by drainage of large lakes. Martian canyons, such as Valles Marineris, may have once contained large lakes, but the source of the sediments, their mode of deposition, and the lifetime of the lakes all remain undetermined. Although the rate of valley formation tailed off at the end of the Noachian, valleys continued to form at a low rate. Their origin could be due to lava flows, maybe caused by large impacts and/or volcanic eruptions. Their presence supports the occasional temporary availability of liquid water on the volcano surfaces.
The main characteristics of each Martian epoch are reported in Table **3.2**. To sum up, strong evidences support the claim that liquid water was present in a great quantity on the martian surface until the Noachian age. At the end of this period, weathering and erosion rates declined rapidly to very low rates, which resulted in dominantly cold surface conditions and development of a thick cryosphere. At this point, liquid water could exist just in particular and isolated areas due to local phenomena, mainly related to volcanic activity.

Table 3.2: Martian geological epochs and their characteristics

Epoch	**Age Range**	**Crater Density**	**Dominant Processes**	**Water-Related Features**
Noachian	4.5–3.7 Ga	> 200 craters / 10^6 km^2 ($D > 5$ km)	Heavy bombardment; valley network formation; phyllosilicate alteration	Extensive valley networks; open-basin lakes; widespread phyllosilicates (clays)
Hesperian	3.7–3.0 Ga	67–200 craters / 10^6 km^2 ($D > 5$ km)	Widespread volcanism; formation of large outflow channels	Large outflow channels; sulfate deposits; chaos terrains
Amazonian	3.0 Ga–present	< 67 craters / 10^6 km^2 ($D > 5$ km)	Impact cratering; aeolian resurfacing; polar ice accumulation	Glacial features; gullies; recurring slope lineae; polar ice caps

Current estimates of the Martian water inventory include: polar ice caps containing approximately $1.6-3.2\times10^6$ km^3 of water ice [27, 111], equivalent to a global layer $11-25$ meters deep; atmospheric water vapor, highly variable seasonally but averaging ~ 10 precipitable microns globally [112]; adsorbed water in the regolith, potentially significant at high latitudes [113]; and subsurface ice deposits detected by radar sounding and neutron spectroscopy, extending to mid-latitudes [114].

3.1.3 Evidence for past surface water

Evidences of abundant liquid water on early Mars rests on multiple, independent lines of evidence spanning geomorphological, mineralogical, sedimentological, and geochemical observations. These data, collected over decades of orbital and in-situ investigations, provide compelling evidence that Mars experienced a hydrologically active past fundamentally different from its current arid state. The most visually evidence for past surface water comes from large-scale fluvial and lacustrine features preserved in the Martian landscape. Valley networks, first identified in Mariner 9 images [115] and subsequently mapped in detail using Mars Global Surveyor, Mars Express, and Mars Reconnaissance Orbiter data, exhibit drainage patterns morphologically analogous to terrestrial river systems.

In addition to that, numerous paleolake basins have been identified through integrated topographic, spectral, and stratigraphic analyses. The dimensions of these basins range from small closed basins a few kilometers across to large open-basin systems exceeding 100 km in diameter [116]. Jezero crater, the Mars 2020 Perseverance rover landing site, represents a well-preserved paleolake basin containing two prominent fan deposits with channels, layered stratigraphy, and spectral signatures of carbonate and hydrated silica-minerals that are indicative of aqueous deposition [117].

Large-scale outflow channels, though representing a different hydrological regime than valley networks, provide evidence for catastrophic water releases from subsurface reservoirs. The presence of features, such as streamlined islands, longitudinal grooves, inner channel terrace, and scoured bedrock are difficult to be explained by any fluid other than water, although the role of CO_2, brines, or ice-water mixtures is still debated [118].

Phyllosilicates are minerals that form in the presence of liquid water. They have been extensively detected on Mars, with most occurrences in Noachian-aged terrains. Phillosilicates can form in several geological contexts: in ancient crater walls and massifs, suggesting early crustal alteration; associated with valley networks and paleolake deposits; in layered sedimentary units within crater basins; and in proximity to volcanic centers, indicating hydrothermal alteration [119]. This diversity suggests multiple aqueous environments active on early Mars, from regional groundwater systems to localized hydrothermal surface water bodies.

Mars rovers have provided ground-truth observations that complement orbital datasets. The Mars Exploration Rover Opportunity, operating in Meridiani Planum from $2004 - 2018$, documented extensive sedimentary outcrops exhibiting: millimeter-scale cross-bedding and structures diagnostic of subaqueous dune migration; spherical hematite concretions formed through groundwater diagenesis; sulfate-rich evaporite minerals formed in acidic, oxidizing aqueous conditions; and erosional surfaces indicating multiple episodes of sediment deposition and erosion [120]. The Curiosity rover in Gale crater has revealed a sedimentary sequence spanning hundreds of meters of stratigraphic thickness, recording the evolution from a fluvio-lacustrine environment to an arid aeolian system [121].

Isotopic measurements provide independent constraints on Mars' water history. The current atmospheric D/H ratio indicates substantial preferential loss of lighter hydrogen isotopes over geological time [122]. This enrichment implies an ancient water inventory significantly larger than the present-day remnant, with estimates suggesting an equivalent global layer of $100 - 1500$ meters if spread uniformly over the planet's surface [123].

Additional isotopic evidence comes from Martian meteorites. These samples reveal: D/H ratios in ancient hydrous minerals lower than present atmospheric values but enriched relative to terrestrial

ocean water, indicating water loss began early in Martian history [124]; oxygen isotope compositions in carbonates and sulfates providing constraints on formation temperatures and water sources [125]; and noble gas isotope ratios indicating extensive atmospheric escape driven by solar wind erosion after magnetic field cessation [126].
To conclude, the convergence of geomorphological, mineralogical, sedimentological, geochemical, and isotopic evidence provides a robust foundation for the conclusion that liquid water was widespread on early Mars, particularly during the Noachian epoch.

3.2 The MARSIS instrument

The Mars Advanced Radar for Subsurface and Ionospheric Sounding (MARSIS) on Mars Express is the first sounding radar operating from orbital altitudes since the Apollo 17 Lunar Sounder flown in 1972 [127]. The MARSIS instrument is a low-frequency, nadir-looking, pulse-limited radar sounder that employs synthetic aperture techniques. Operating on-board Mars Express spacecraft, which follows a highly elliptical 7−hour orbit around Mars, MARSIS conducts subsurface sounding passes lasting $26 - 30$ minutes. During these operations, the instrument shares spacecraft power and data resources with five other instruments aboard Mars Express. The observation frequencies, which will be introduced in subsequent sections, were selected as a compromise between maximizing signal penetration into the subsurface and ensuring effective propagation through the Martian ionosphere.

3.2.1 Instrument characteristics and orbital configuration

MARSIS operates as a synthetic aperture radar sounder transmitting low-frequency chirped pulses and receiving echoes from both the surface and subsurface interfaces characterized by dielectric contrast. As it will be shown in detail in Section **3.2.2**, MARSIS can operate in two different modes, according to the scientific goal of the observation: subsurface and ionospheric. The instrument consist of three principal subsystems that work together to enable both ionospheric and subsurface sounding mode capabilities:

- the antenna subsystem comprises a 40−meter dipole antenna and a 7−meter monopole antenna. The dipole antenna operates in dual mode, for transmission and reception of sounder pulses, while the monopole antenna operates in receive-only mode specifically designed to detect surface clutter echoes. The 7−meter monopole antenna is positioned at the center of the dipole, oriented normal to it, in order to minimize coupling between the dipole and monopole while creating a null in the nadir direction. However, during the testing phase of the instrument on ground, the noise level in the monopole channel proved higher than expected, significantly limiting its utility to only a few data acquisitions [39].
  The dipole antenna pattern at MARSIS frequencies exhibits broad beamwidths (approximately $60 - 90°$ at half power) that result in large footprints on the surface. The radius of the illuminated area (Fresnel zone, Section **1.4.4**) at nadir is given by Equation **1.59**. For typical values of MARSIS subsurface observations, such as $h = 800$ km and $\lambda = 100$ m (Band 2), the corresponding Fresnel zone radius is ~ 6.3 km, producing an illuminated area of ~ 125 km^2. This large footprint shows why SAR processing is required to achieve adequate along-track resolution.
  In the receiving mode, the thermal noise floor of MARSIS is circa -115 dBm (referenced

to 1 MHz bandwidth). Surface echoes from nadir typically produce received powers of -70 to -90 dBm, providing signal-to-noise ratios of $25 - 45$ dB for surface returns. Subsurface echoes, attenuated by two-way propagation through the overlying material and reflection coefficient losses at the interface, are significantly weaker: typically $20 - 40$ dB below the surface return for ice-rock interfaces, placing them within $5 - 15$ dB of the noise floor. This necessitates coherent integration of multiple pulses along-track to improve SNR, a process that requires careful management of phase coherence as analyzed in this work.

- The Radio Frequency Subsystem encompasses both transmit and receive channels for the dipole and monopole antennas. The transmitter consists of two separate modules: one covering Band 1 ($1.3-2.3$ MHz) and another covering Bands 2, 3, and 4 ($2.5-5.5$ MHz). The frequency gap between 2.3 and 2.5 MHz provides improved isolation between transmitters. The receiver must meet extremely demanding performance requirements. It must exhibit very low sidelobes in its impulse response to detect weak subsurface returns in the presence of strong surface returns, while maintaining high linearity to prevent unwanted harmonic generation. The received signal passes are processed as follows: initial RF amplification, selectable bandpass filtering (with Band 1 having larger bandwidth to accommodate ionospheric sounding mode), downconversion with local oscillator mixing, low-pass filtering, further amplification with gain control, and finally analog-to-digital conversion using two $8-$bit converters sampling at 2.8 MHz.
- The Digital Electronics Subsystem (DES) implements all instrument logic, most spacecraft interfaces, and the onboard data processing capability. It is composed of several different components, among which: a reference oscillator providing instrument coherence, a sounder signal generator, a timing generator, an echo processor, and dedicated power converters. In particular, the reference oscillator is a numerically controlled oscillator clocked at 28 MHz. The sounder signal generator is a digital chirp generator capable of producing both linear frequency-swept signals and pulsed monochromatic signals depending on the operating mode.

The functional block diagram of MARSIS is presented in Figure **3.1**.
Mars Express operates in a highly elliptical polar orbit with its approximate parameters reported in Table **3.3**. This eccentric orbit was optimized for both high-resolution imaging at low altitudes and global coverage. However, for MARSIS subsurface sounding operations, the spacecraft must be at altitudes exceeding ~ 240 km to avoid interference between surface and ionospheric echoes and to

Table 3.3: Orbital parameters of Mars Express

Parameter	**Value**
Periapsis altitude	250–350 km
Apoapsis altitude	10,100–11,500 km
Orbital period	∼7 hours
Inclination	∼86.9° (near-polar)
Periapsis latitude	Variable (shifted over mission lifetime)

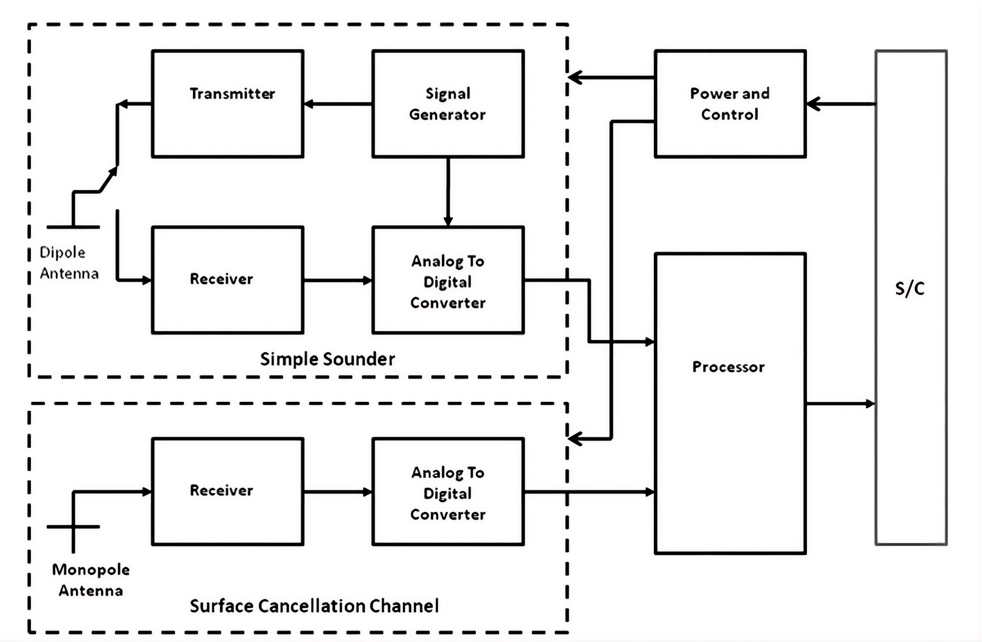

Figure 3.1: Conceptual block diagram of MARSIS. From [128].

ensure safe antenna clearance. Consequently, MARSIS operates primarily during orbital segments at altitudes between 240 − 900 km, providing optimal geometry for subsurface investigations while maintaining sufficient signal-to-noise ratio [39].

The spacecraft's ground track velocity at typical MARSIS operating altitudes is approximately:

$$v_{\text{ground}} \approx 3.3\,\text{km/s} \tag{3.1}$$

At the nominal PRF of 127 Hz, this corresponds to an along-track pulse spacing of:

$$\Delta x = \frac{v_{\text{ground}}}{\text{PRF}} \approx 26\,\text{m} \tag{3.2}$$

This dense spatial sampling enables coherent integration and SAR focusing to improve the along-track resolution beyond the antenna beamwidth limitation.

3.2.2 Operating modes and frequency bands

MARSIS operates in multiple distinct modes, each optimized for specific scientific objectives and adapted to different orbital geometries and ionospheric conditions. The mode selection and frequency band allocation reflect compromises between penetration depth, resolution, ionospheric propagation effects, and data volume constraints.

Subsurface Sounding (SS) Modes

Subsurface sounding operates over a 1.3 − 5.5 MHz frequency range across four frequency bands, each with 1 MHz bandwidth. The subsurface sounder employs two separate channels: the first connected to the dipole antenna containing both transmitter and receiver, and the second connected to the receive-only monopole antenna. The monopole antenna's design, as already mentioned, with its null in the nadir direction, is specifically intended to receive only returns from the surface away from nadir, helping discriminate between top surface and subsurface returns.

The choice of operating frequency for a given observation involves trade-offs between several factors.

Penetration depth decreases with increasing frequency due to frequency-dependent attenuation in geological materials. From Section **1.3**, the skin depth δ can be rewritten as:

$$\delta = \frac{c}{2\pi f \sqrt{\epsilon_r} \tan \delta_e} \tag{3.3}$$

Where ϵ_r is the relative permittivity and tan δ_e is the loss tangent. For water ice at Martian polar temperatures ($\sim 150 - 200$ K), the loss tangent is approximately 10^{-4} to 10^{-3}, yielding to one-way attenuation rates of $0.5 - 2$ dB/km at 4 MHz [129]. On the other hand, dry basaltic materials exhibit loss tangents of 10^{-2}, producing attenuation rates of $5 - 50$ dB/km, preventing significant penetration [5].

The ionosphere, produced by solar UV ionization of the predominantly CO_2 atmosphere, introduces both attenuation and phase dispersion that increase with decreasing frequency. Total electron content (TEC) along the propagation path can reach $10^{16} - 10^{17}$ electrons/m^2 during solar maximum at solar zenith angles near zero, causing significant signal degradation at the lowest MARSIS frequencies [130]. The critical frequency f_c of the Martian ionosphere, that is, the frequency below which vertical propagation is impossible due to total reflection, varies from ~ 0.5 MHz at night to $2 - 4$ MHz during daytime at solar maximum [131]. Operations at 1.8 MHz are therefore possible only at high solar zenith angles or during solar minimum conditions, while 3.0 and 4.0 MHz bands provide more reliable ionospheric penetration. This ionospheric interaction represents one of the primary challenges in MARSIS data processing, as discussed in subsequent sections.

Surface clutter from off-nadir surface returns interferes with subsurface echo detection. Higher frequencies, with shorter wavelengths and thus smaller Fresnel zones, reduce the spatial extent of surface clutter contamination. This effect partially neglects the penetration advantage of lower frequencies [132].

The instrument offers five distinct subsurface sounding modes (SS1-SS5), each optimized for different scientific objectives and operational constraints [128]:

- SS1 operates with 2 frequency bands, 2 antennas, and 1 Doppler filter, allowing coherent clutter cancellation on two frequency bands through dual antenna clutter cancellation in ground processing.

- SS2 uses 2 frequency bands, 1 antenna (dipole only), with onboard multi-look processing that performs non-coherent integration of five looks and downloads a single amplitude-detected averaged echo profile for each synthetic aperture.

- SS3 employs 2 frequency bands, 1 antenna, and 3 Doppler filters, downloading I/Q data for three Doppler filters collected on the dipole antenna channel at two frequencies, with range processing performed on the ground.

- SS4 operates with 1 frequency band, 2 antennas, and 5 Doppler filters, enabling ground-based dual antenna cross-track clutter cancellation with five-look non-coherent integration.

- SS5 utilizes 1 frequency band, 2 antennas, a short pulse (30 microseconds), and 3 Doppler filters. This mode uses a short-pulse waveform to reduce the impact of uncontrolled sidelobes on deep subsurface reflections.

The complete list of MARSIS subsurface sounding mode parameters is reported in Table **3.4**.

Table 3.4: MARSIS subsurface sounder parameters

Parameter	**Value**
Subsurface sounding altitude range	250–900 km
Dipole antenna length	40 m tip to tip
Monopole antenna length	7 m
Subsurface sounder frequency range	1.3–5.5 MHz (4 bands)
Peak radiated power	1.5–5 W
Subsurface sounder pulse repetition rate	127 Hz
Subsurface sounder transmit pulse length	250 μs
Subsurface sounder wavelength	60–160 m
Bandwidth per band	1 MHz
Sounder free space depth resolution	150 m
Sounder dynamic range	40–50 dB
Nominal depth window	15 km
Number of processed channels	2 or 4
Number of simultaneous frequencies	1 or 2
Data quantization	8 bits per sample
Data quantization rate	2.8 megasamples per second
Data window duration	350 μs
Data rate output	18–75 Kbps
Data volume	285 Mbit per day
DC operating power	60 W
Total mass	20 kg

Active Ionospheric Sounding (AIS) Mode

The ionospheric sounder operates over a much broader frequency range of 100 kHz to 5.5 MHz using a simple low-bandwidth pulse, radiating only through the dipole antenna. This mode operates when the spacecraft is at altitudes higher than 900 kilometers or when no subsurface mode is being used. AIS mode is designed to characterize the vertical structure of the Martian ionosphere. The pulses are transmitted in sequence, spanning the entire bandwidth, at steps of 10.9 kHz, covering 488 discrete frequencies [133]. The operational principle exploits ionospheric reflection: radio waves transmitted at frequencies below the local plasma frequency are reflected back to the spacecraft, with reflection altitude decreasing as transmission frequency increases.
The plasma frequency f_p can be written as:

$$f_p = \frac{1}{2\pi}\sqrt{\frac{n_e e^2}{m_e \epsilon_0}} \approx 9\sqrt{n_e}\ \text{Hz} \tag{3.4}$$

Where e is the electron charge, m_e the electron mass, and n_e is the electron density expressed in electrons/m^3. By measuring the time delay to the reflected signal as a function of transmitted frequency, the altitude profile of electron density can be inverted to produce vertical ionograms [134].
The AIS mode can operate only when no SS mode is active. During several orbital passes, MARSIS alternates between SS and AIS modes to simultaneously characterize both the subsurface and overlying ionosphere. Usually, a typical measurement cycle consists of an AIS sweep lasting ~ 90 seconds, followed by SS observations at selected frequency band for several minutes. This approach allows direct measurement of ionospheric TEC at the same time of subsurface sounding observations, enabling the estimation of a proper ionospheric correction. The downside is that this procedure creates gaps in spatial coverage that can complicate SAR focusing for subsurface targets.
The complete list of MARSIS ionospheric sounding mode parameters is reported in Table **3.5**.

Table 3.5: MARSIS ionospheric sounder parameters

Parameter	**Value**
Ionospheric sounding altitude range	∼1200–250 km
Ionospheric sounder frequency range	100 kHz –5.5 MHz
Discrete frequencies used for sounding	160
Ionospheric sounder transmit pulse length	91.4 μs
Ionospheric sounder pulse repetition rate	127 Hz
Receive window after pulse	254–7566 μs
Ionogram frame length	1.26 s
Ionogram frame interval	7.543 s
Detection bandwidth	10 kHz

3.2.3 Data acquisition geometry

Understanding the acquisition geometry of MARSIS observations is crucial for interpreting radar data and developing appropriate processing algorithms. The geometry takes into account complex interactions between spacecraft motion, antenna radiation patterns, surface/subsurface topography, and electromagnetic wave propagation through the ionosphere and subsurface.
The coordinate systems that are involved in a MARSIS observation are:

- Mars-centered inertial frame (MCI): Fixed reference frame with origin at Mars' center of mass, used for spacecraft trajectory specification.
- Mars-centered, Mars-fixed frame (MCMF): Rotating frame fixed to Mars' surface, with orientation defined by IAU Mars coordinate conventions. Surface features are referenced in this frame using areocentric latitude (ϕ), east longitude (λ), and radius from Mars center (r).
- Spacecraft-centered frame: Origin at spacecraft center of mass, with axes typically defined by velocity vector (along-track), normal to orbital plane (cross-track), and nadir direction (radial).
- Antenna frame: Dipole antenna oriented perpendicular to spacecraft-Sun direction, defining the polarization of transmitted and received electromagnetic waves.

A typical MARSIS observation of Mars consists of a sequence of Frames. Frames contain a variable number of pulses, as a function of altitude and frequency. The operational architecture is designed to optimize data quality while managing spacecraft resource constraints (data rate, power, telemetry volume). A sketch of the frame structure is shown in Figure **3.2**. Each Mars observation Frame is made of the following sequence of operations performed on-board [135]:

- initial orbital parameters estimation, including Frame size estimation (NB, number of PRIs) and synthetic aperture size (NA);
- signal transmission (2 pulses) and echoes reception, repeated NA times;
- NA echoes coherent weighted accumulation (Synthetic Aperture formation);
- Doppler processing to synthesize three Doppler filters;
- range compression on the central Doppler filter;
- surface echo tracking to update reception parameters for subsequent Frame.

Frame size NB is computed adaptively during each observation to ensure that contiguous synthetic apertures are properly separated. The spacecraft displacement during NB pulses is:

$$\Delta S = \sqrt{\frac{\lambda_1 \cdot H}{2}} + N_0 \cdot \frac{V_{\tan}}{\mathrm{PRF}} \tag{3.5}$$

where N_0 is a constant offset of 36 PRIs, λ_1 is the wavelength of the lowest operative frequency in use, H and $V_{\tan}$ are the spacecraft altitude and the tangential velocity (Equation **3.1**).
Frame size is then computed as:

$$\mathrm{NB} = \mathrm{Int}\left[\frac{\Delta S}{V_{\tan}} \cdot \mathrm{PRF}\right] \tag{3.6}$$

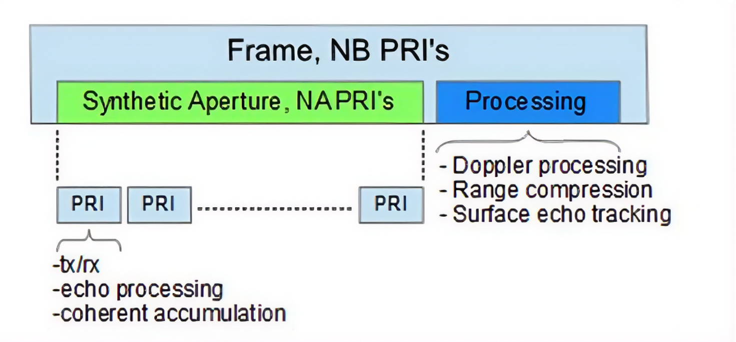

Figure 3.2: MARSIS frame structure. From [135].

This adaptive sizing guarantees that the spacing between consecutive synthetic apertures precisely matches the spacecraft travel distance, maintaining optimal coverage without gaps or overlaps. Synthetic aperture sizes NA_1 and NA_2 are computed separately for the two operative frequencies in dual-frequency mode:

$$\mathrm{NA}_i = \mathrm{Int}\left[\lambda_i \cdot \frac{H \cdot \mathrm{PRF}}{2 \cdot \gamma_i \cdot V_{\tan} \cdot \Delta S}\right] \qquad i = 1, 2 \tag{3.7}$$

where γ_1 and γ_2 are corrective frequency-dependent values necessary to obtain the same azimuth resolution at different frequencies. The actual synthetic aperture size used is $\mathrm{NA} = \max(\mathrm{NA}_1, \mathrm{NA}_2)$, ensuring sufficient aperture length for both frequency channels.

MARSIS transmits two square pulses each of width 250 μs modulated in frequency with a 1 MHz bandwidth centered on the selected Operative Frequency. The time delay between the two transmitted pulses is fixed at 450 μs. The Trigger value, which positions in time the receive window to detect surface and subsurface echoes, is adaptively computed for each Frame accounting for spacecraft altitude and ionospheric delay. For the first Frame (Frame 0), the Trigger is calculated as:

$$\mathrm{Trigger} = \frac{2H}{c} + \Delta t \tag{3.8}$$

where Δt is a preset offset compensating for ionospheric delay. For subsequent Frames ($n \geq 1$), Trigger values are estimated using results from surface echo tracking executed on the previous Frame (Frame $n-1$), enabling adaptive tracking of surface topography variations along the orbital track.

During each PRI, received echoes are sampled at 980 samples per echo, with $8-$bit quantization. Subsequently, during the synthetic aperture accumulation, MARSIS performs coherent weighted accumulation of NA echoes, applying different phase factors to synthesize three Doppler filters with distinct pointing geometries:

- Doppler Filter 0: Nadir-pointing;
- Doppler Filter -1: Forward-looking (ahead of spacecraft);
- Doppler Filter +1: Backward-looking (behind spacecraft).

This multi-filter approach provides spatial diversity, enabling observation of different surface areas simultaneously and facilitating off-nadir clutter characterization. The central filter (Filter 0) undergoes range compression on-board, while all three filters are transmitted to ground for further processing and analysis.

In addition to the standard sequential Frame architecture, MARSIS can operate in Super-Frame mode (Figure **3.3**), where a single, extended Frame is executed instead of multiple short Frames. This operational mode is achieved by increasing the offset parameter N_0 in the Frame size calculation (Equation **3.2**), forcing the instrument to treat an extended observation interval as a single coherent aperture. Super-Frame mode offers several advantages for subsurface sounding:

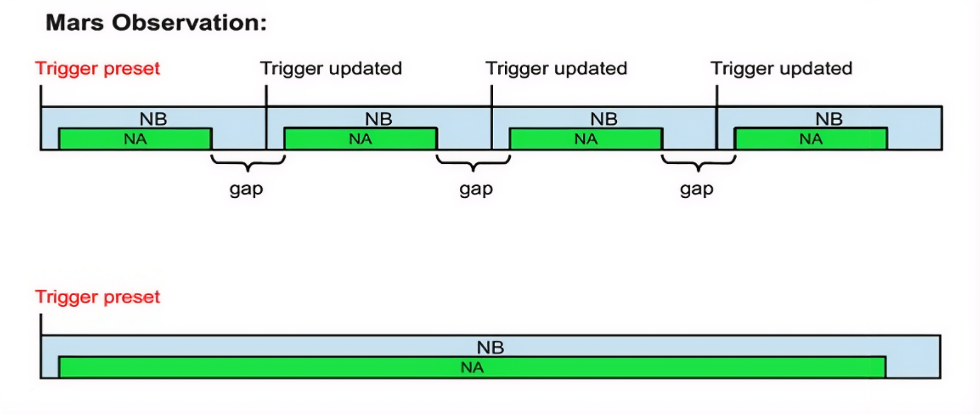

Figure 3.3: MARSIS Frame and Super-Frame modes. From [135].

- by maintaining phase coherence over longer synthetic apertures (up to 3200 PRIs), Super-Frames enable improved along-track resolution and signal-to-noise ratio for weak subsurface returns.
- Super-Frame operations utilize the on-board Flash Memory (16.77 MByte capacity) to store unprocessed radar echoes. Each PRI generates 980 samples × 2 channels × 1 byte = 1960 bytes, allowing storage of:

  $$N_{\max} = \frac{16.77 \times 10^6 \,\text{bytes}}{3200\,\text{PRI} \times 2\,\text{channels} \times 980\,\text{samples} \times 1\,\text{byte}} \approx 2.6 \text{ Super-Frames} \tag{3.9}$$

  the duration of a single Super-Frame is:

  $$T_{SF} = \frac{\text{NA}}{\text{PRF}} \approx 25.1\,\text{seconds} \tag{3.10}$$

  This raw data can be transmitted and processed on ground.
- Standard Frame processing introduces discontinuities at Frame boundaries where synthetic aperture integration resets. Super-Frames eliminate these artifacts, providing uniform processing across the entire observation window.
- Raw Super-Frame data enables ground-based implementation of advanced focusing algorithms, ionospheric correction methods, and clutter suppression techniques that require access to unfocused, uncorrected echo data.

After acquiring a Super-Frame, the instrument requires ~ 44 seconds to transfer data from RAM buffers to Flash Memory devices, during which no new data acquisition can occur. This duty cycle limits continuous coverage but is acceptable for targeted observations of specific regions of interest. A typical MARSIS subsurface sounding observation operates continuously for up to ~ 30 minutes, covering a ground track approximately 1200 km in length. At the along-track sampling of ~ 26 m, this corresponds to approximately:

$$N_{\text{total}} = \frac{1200}{26} \approx 46000 \text{ along-track samples} \tag{3.11}$$

organized into sequential Frames, each containing one synthetic aperture.

3.3 The study area: Ultimi Scopuli region

3.3.1 The South Polar Layered Deposit (SPLD)

The South Polar Layered Deposit (SPLD, Figure **3.4**) of Mars consists in one of the planet's most consistent ice accumulations, containing an estimated 1.6 million km^3 of material, equivalent to a global layer of water ~ 11 meters deep if uniformly distributed over the planet's surface [111]. These deposits, centered on Planum Australe and extending to approximately 75°S latitude, show a distinctive dome-shaped morphology rising up to $3 - 4$ km above the surrounding plains, with a maximum thickness exceeding 3.7 km in some locations.

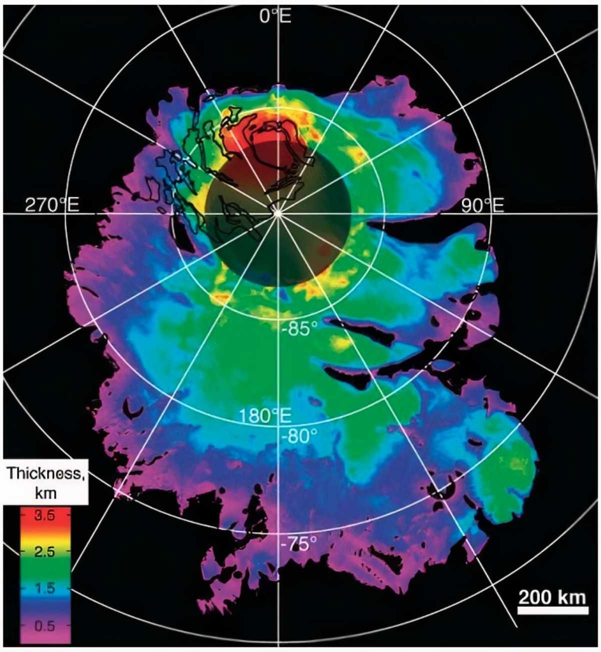

Figure 3.4: Map of the SPLD thickness, based on MARSIS measurements and MOLA surface topography. From [111].

The SPLD are characterized by remarkably regular internal stratification visible in orbital radar sounding data, consisting of alternating layers with varying dielectric properties that reflect compositional or physical property variations. This layering, with typical layer thicknesses ranging from a few meters to several tens of meters, is interpreted as a record of climate variations driven by periodic changes in Mars' orbital parameters (obliquity, eccentricity, and precession) over timescales of $10^4 - 10^6$ years [136].

The SPLD are primarily composed of water ice mixed with varying amounts of atmospheric dust. Density measurements derived from gravity and topography data yield bulk densities of $1,100 - 1,400$ kg/m^3 [137, 138], intermediate between pure water ice (~ 920 kg/m^3) and typical Martian basaltic dust ($\sim 2,900$ kg/m^3). These density constraints indicate dust mass fractions ranging from approximately 2% to 20%, with spatial variations reflecting temporal changes in atmospheric dust loading and ice accumulation efficiency [138]. The dielectric properties of the SPLD, crucial for interpreting radar sounding data, are determined by the ice-dust mixture composition and temperature. As mentioned in Section **3.2.2**, pure water ice at Martian polar temperatures exhibits a real relative permittivity of approximately 3.1-3.15 and a low loss tangent of 10^{-4} to 10^{-3} at MARSIS frequencies, resulting in minimal electromagnetic attenuation [5]. The incorporation of dust increases both the permittivity and loss tangent: for a 10% dust content (by mass), the effective permittivity increases to $\sim 3.5 - 4.0$, while attenuation rates remain sufficiently low ($0.5 - 2$ dB/km) to permit radar penetration to depths of several kilometers [139].

Temperature exerts a strong influence on ice loss tangent at MARSIS frequencies, with warmer ice exhibiting significantly higher attenuation than cold ice [140]. Surface temperatures of the SPLD, measured by thermal infrared spectroscopy, average $\sim$ 160 K, while basal temperatures depend on geothermal heat flux, thermal conductivity of the overlying material, and potential heating from viscous deformation or basal melting.

The uppermost surface of the SPLD in the study region is smooth at MARSIS wavelengths (75-167 m). Topographic data from the Mars Orbiter Laser Altimeter (MOLA) reveal surface roughness well below the MARSIS wavelength, indicating that scattering is predominantly coherent (specular reflection) rather than diffuse [141]. This smoothness results in surface echoes originating primarily from the nadir direction, simplifying radar data interpretation by minimizing off-nadir surface clutter contamination.

Observations from the High Energy Neutron Detector (HEND) on Mars Odyssey show the presence of a seasonal layer of CO_2 ice, less than 1 meter thick, over most of Planum Australe. This layer is present during winter, while it sublimates during spring and summer [142]. At MARSIS frequencies and for thicknesses <1 m, this seasonal CO_2 layer has negligible impact on surface reflectivity or subsurface echo characteristics [143].

The internal structure of the region has been characterized by MARSIS and Shallow Radar (SHARAD) radar sounding observations. At MARSIS frequencies the penetration depth is larger, allowing the detection of both surface returns and echoes from the base of the deposit or from internal dielectric discontinuities, when present. SHARAD operates at higher frequencies (20 MHz) with a larger bandwidth (10 MHz) that improves significantly the range resolution (15 m), revealing fine-scale internal layering throughout most of the SPLD. However, sometimes it does not detect basal echoes due to cumulative signal attenuation [132].

The internal layering visible in radar data exhibits lateral continuity extending over hundreds of kilometers, with layer-parallel geometry indicating relatively uniform accumulation and minimal post-depositional deformation in most regions. According to [144], four distinct radar facies have been identified by SHARAD within SPLD, reflecting variations in layer properties and depositional history:

- Shallow Layers (SL): Closely spaced, well-focused near-horizontal reflectors in the uppermost stratigraphic position. Average thickness $\sim$ 54 m, covering $\sim$ 14% of the SPLD area, concentrated between 270°E and 90° E near the residual cap.
- Focused Layers (FL): Well-resolved, distinct reflectors (1-2 range cells wide) extending over much of the SPLD depth, subdivided into normal brightness (FLN) and substantially brighter reflectors (FLB).

- Blurred Layers (BL): Evident reflectors that are vertically unfocused (> 2 range cells wide, up to 6 range cells). Cover $\sim 12\%$ of SPLD area, concentrated in the thickest regions of the deposits. They represent a continuum with FL, suggesting similar stratigraphy with different radar properties.
- Low Reflectivity Zones (LRZ): Regions with few or no distinct reflectors, covering $\sim 49\%$ of the SPLD surface area—the most extensive facies. Occurrences near the South Pole identified them as CO_2 ice deposits [145], though composition of other LRZ remains uncertain. The largest continuous LRZ deposit ($131 - 255°$E) varies in thickness from ~ 370 m (west) to ~ 80 m (east).

The Ultimi Scopuli study region (Figure **3.5**), located at approximately 81° S, 193° E, falls within an area characterized by LRZ in SHARAD data, where internal layering is barely discernible, which corresponds to relatively continuous layering in MARSIS data.

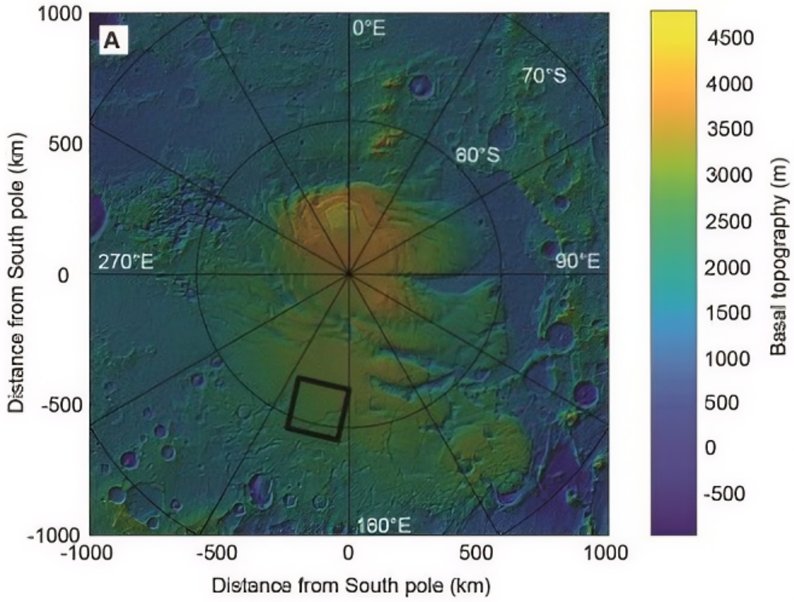

Figure 3.5: Shaded relief map of Planum Australe, Mars, south of 75°S latitude. The map was produced using the MOLA topographic dataset. The black square outlines the study area. From [143].

3.3.2 The bright basal reflection

In 2018, [143] reported the detection of an anomalously bright basal reflection in MARSIS radar data from a localized region of the Ultimi Scopuli area in the South Polar Layered Deposits. This discovery, based on analysis of 29 orbital observations acquired between May 2012 and December 2015, identified a ~ 20 km wide zone centered at approximately 81° S, 193° E where subsurface echoes exhibited characteristics inconsistent with a simple ice-rock interface and potentially indicative of liquid water beneath the ice cap. The bright basal reflector extends laterally for 20 km, forming a roughly elliptical area embedded within the broader SPLD. This feature has been consistently observed at all three MARSIS frequencies employed over the study area (3,

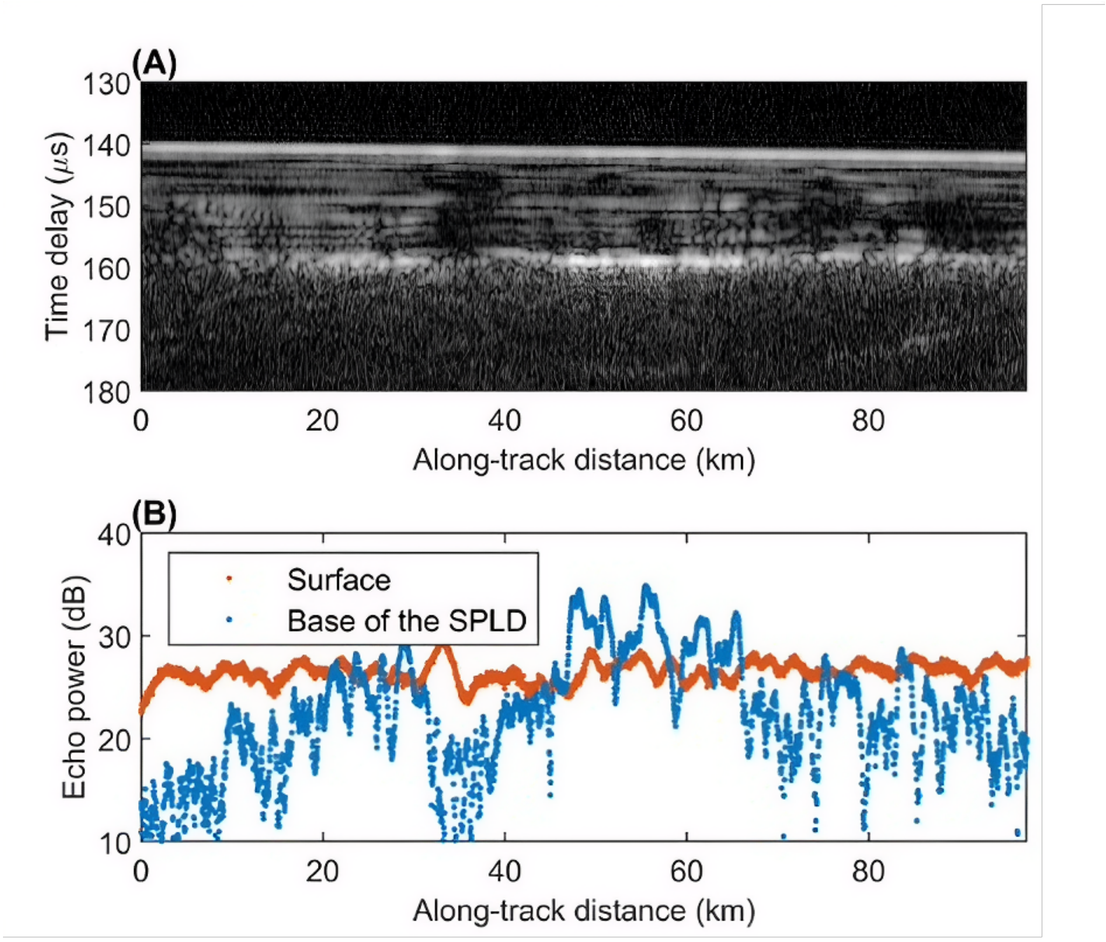

Figure 3.6: **A**: Radargram for MARSIS orbit 10737. **B**: Plot of surface and basal echo power for the radargram in **A**. From [143].

4, and 5 MHz), with similar spatial distribution and relative intensity at each frequency. This multi-frequency detection provides confidence that the feature is not an artifact of ionospheric distortion or instrument effects, which would be expected to vary differently with frequency.

Figure **3.6** shows the radargram of the orbit 10737, which is presented in the aforementioned work and it constitutes the major component of this analysis. This observation has been acquired in Super-Frame mode, collecting together 3200 echoes that have been processed on ground. In the radargram, the surface return appears as a continuous bright line, while beneath it, internal layering of the SPLD is faintly visible. At approximately 160 μs two-way travel time, a basal reflection appears, exhibiting unusual characteristics: between 45 and 65 km along-track, this basal echo becomes dramatically brighter, locally exceeding the surface echo power by several decibels even after accounting for two-way attenuation through 1.5 km of ice.

In typical radar sounding scenarios, subsurface echoes are invariably weaker than surface returns due to transmission losses at the surface, geometric spreading losses, attenuation within the intervening material, and typically lower reflection coefficients at subsurface interfaces compared to the air-ice surface boundary. Quantitatively, for an orbital radar sounder illuminating a smooth, planar surface, the received surface echo power is:

$$P_s = P_t \times \left(\frac{G\lambda}{8\pi H}\right)^2 \times R_s \tag{3.12}$$

being P_s and P_t the received and transmitted power, respectively, and R_s is the reflection coefficient defined in Equation **1.46**.

For a subsurface interface at depth z below the surface, the received echo power is:

$$P_{ss} = P_t \times \left(\frac{G\lambda}{8\pi(H+z)}\right)^2 \times R_{ss} \times T_s^2 \times \exp(-2\pi f \tan(\delta)\,\tau) \tag{3.13}$$

where T_s is the power transmission coefficient at the surface, R_{ss} is the reflection coefficient at the subsurface interface, tan (δ) is the loss tangent of the ice, and τ is the two-way travel time through the ice to the interface. The exponential term quantifies signal attenuation due to dielectric losses during two-way propagation. Since the antenna was not calibrated on ground, one common approach for quantitative data analysis is based on the ratio of subsurface to surface echo power::

$$\frac{P_{ss}}{P_s} \approx \frac{R_{ss}}{R_s} \times T_s^2 \times \exp(-2\pi f \tan\delta\,\tau) \tag{3.14}$$

where the approximation holds under the assumption that the spacecraft altitude is much larger than the subsurface depth, so that the geometric spreading terms cancel. This ratio isolates the effects of basal interface properties and attenuation from instrumental factors like transmitted power and antenna gain. Moreover, it makes pointless the application of ionospheric correction algorithms, since phase distortion introduced by ionospheric delay impacts surface and subsurface in the same way. Also, given typical MARSIS altitudes of 300-800 km versus subsurface depths of $\sim$ 1.5 km: $H \approx H + z$. For the bright reflector region, measured values of P_{ss}/P_s reach 0-3 dB (ratios of 1-2) at different frequencies, indicating that the subsurface echo is comparable to or stronger than the surface echo. Outside the bright area, typical values are -15 to -5 dB (ratios of 0.03-0.3), more consistent with normal subsurface reflections.

The ground processing strategy was tailored for this particular scenario. Topographic data from MOLA reveal exceptional smoothness at MARSIS wavelengths: under these conditions, surface scattering is almost entirely coherent (specular reflection from nadir), and SAR azimuth focusing provides no resolution improvement: this is the reason why the analysis has been conducted over range-compressed only data.

The normalized basal echo power, once corrected for geometry and surface effects, provides a diagnostic for basal interface permittivity. To extract this value from measured P_{ss}/P_s ratio, [143] developed a one-dimensional electromagnetic plane wave propagation model to simulate MARSIS echoes for a three-layer stratigraphy (Section **3.4.2**): free space above the spacecraft, a homogeneous SPLD layer representing the ice sheet, and a semi-infinite basal substrate with unknown permittivity. For each combination of parameters, the model computed frequency-dependent surface and basal reflection coefficients, propagation losses through the ice, and resulting P_{ss}/P_s values at MARSIS frequencies (3, 4, 5 MHz). This generated a family of curves relating normalized basal echo power to basal permittivity for different SPLD dust contents and temperatures. The measured distributions of normalized basal echo power inside and outside the bright reflector region were then inverted using these model curves. For each data point, the range of basal permittivity values consistent with the observed P_{ss}/P_s was determined by intersecting the measurement with the family of model curves spanning the parameter space. This yields not a single permittivity estimate but a probability distribution reflecting uncertainties in SPLD composition and temperature.

The estimated basal permittivity values within the bright reflector have a median $\sim 15 - 20$, with significant probability mass extending to 30-40 and beyond, and it cannot be explained by any plausible solid material. The permittivity range 15-40 falls between pure liquid water and concentrated aqueous brines. For Mars, perchlorate brines are of particular interest: magnesium and calcium perchlorate solutions can remain liquid at temperatures as low as 200-210 K for salt concentrations of ∼ 30 weight percent [146], well within the plausible range for basal temperatures given geothermal heating. The localized nature of the bright reflector (a ∼ 20 km diameter region embedded within otherwise normal reflectivity soil) suggests a confined body of liquid rather than a continuous basal aquifer, a scenario that has analogous to terrestrial subglacial lakes beneath Antarctic and Greenland ice sheets.

3.3.3 Scientific context and open questions

The detection of a bright basal reflector beneath the Martian south polar ice cap represents a watershed moment in Mars exploration, potentially identifying the first known body of liquid water on present-day Mars. The implications extend far beyond planetary geophysics, touching fundamental questions in astrobiology, planetary habitability, and future human exploration.

Following the initial 2018 discovery, subsequent analysis of MARSIS data over the south polar region has yielded additional detections that both validate and complicate the interpretation. [1] applied signal processing methods adapted from terrestrial radio-echo sounding studies to discriminate wet from dry subglacial substrates, and analyzed an expanded dataset including observations beyond those in the original study. Their analysis confirmed the initial bright reflector detection and identified additional, smaller areas of anomalously high basal reflectivity in the vicinity of the main feature (see Figure **3.7**), located at similar depths (∼ 1.5 km) but spatially separated by several to tens of kilometers. The identification of multiple bright reflectors close to the main one pose new questions about their nature. Specifically, if they could be isolated liquid body or if they represent different aspects of a single, geometrical complex water reservoir. Their separation and their evolution through time currently represent unresolved questions.

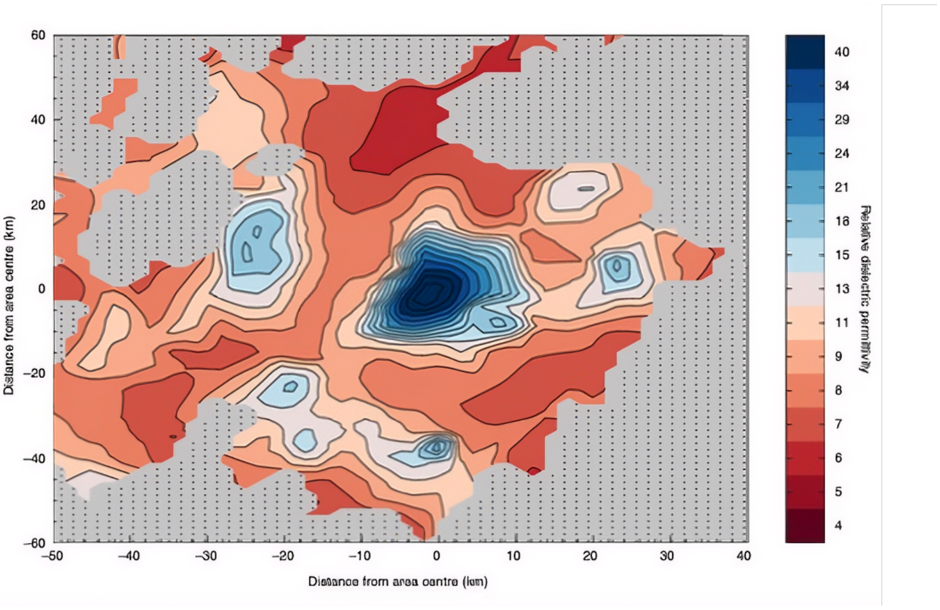

Figure 3.7: Relative dielectric permittivity map computed by inverting the radar data. Values larger than 15 suggest the presence of liquid water. From [1].

The entire analysis we have summarized so far is based on two-dimensional radargrams, that are vertical slices through the subsurface along individual orbital ground tracks. While these 2D profiles provide sufficient information to detect the presence of bright basal reflectors and to estimate their approximate lateral extent (by mapping which orbital passes intersect the region), they offer only limited insight into the three-dimensional structure and spatial relationships that are critical for understanding the features' origin and evolution. By realizing a 3D reconstruction of the area, in principle it is possible to address the following questions:

- the actual three-dimensional shape of the basal interface in the bright reflector region;
- the bright reflector area, and the spatial variation of the basal permittivity;
- the relation with bedrock tomography. Specifically, if the bright reflector(s) occupy topographic lows, slope breaks, or other geomorphologically significant settings;
- the total volume of the high permittivity material.

In the following Sections we will address in detail the challenges of such a 3D reconstruction, given the scenario and the data at disposal. We will show why conventional approaches to three-dimensional reconstruction from radar data can't be applied to this problem and we will present the algorithm developed during this work.
Before describing this reconstruction methodology, it is necessary to address the unique challenges presented by MARSIS data quality and the processing steps required to prepare raw radar echoes for coherent three-dimensional integration. These challenges—ionospheric distortion, residual phase errors, noise characteristics, and spacecraft motion effects—are the subject of the following sections. We specify that Orbit 10737, illustrated in Figure **3.6** and providing one of the clearest expressions of the bright basal reflection, serves as a primary dataset for algorithm development and validation in this work.

3.4 MARSIS data processing pipeline

The processing pipeline described in this section transforms raw Super-Frame radar echoes into range-compressed, topographically aligned radargrams suitable for subsequent analysis (L1C product). Because the observations used in this work were acquired in Super-Frame mode, all 6400 raw echoes (3200 per frequency channel) are available as unprocessed time-domain samples prior to any on-board synthetic aperture formation, allowing full control over every processing step.

3.4.1 Raw data to range compression

The raw data are organized in a matrix of dimensions $N_s \times 2N_x$, where $N_s = 980$ is the number of range samples per echo and $N_x = 3200$ is the number of pulses per frequency channel. The two frequency channels (centered at 3 and 4 MHz or 4 and 5 MHz depending on the operative configuration) are stored in contiguous column blocks and extracted separately. For each pulse, ancillary geometric parameters are read from the geometry table: spacecraft altitude H, radial velocity, sub-spacecraft latitude and longitude. The along-track ground distance is computed from the geodetic coordinates as the cumulative length along the spacecraft ground track:

$$x_k = \sum_{i=1}^{k} \sqrt{(\Delta x_i)^2 + (\Delta y_i)^2 + (\Delta z_i)^2} \tag{3.15}$$

where (x_i, y_i, z_i) are Cartesian coordinates computed from the spherical Mars reference radius $R = 3389.95$ km. This quantity serves as the along-track axis in all subsequent radargrams.

Raw echo amplitudes are stored as 8-bit integers after on-board automatic gain control (AGC) attenuation. The AGC level for each frame is encoded in a specific field, with separate values for the two frequency channels packed into 8-bit words. The actual attenuation in decibels is decoded from the relevant bit fields, and the linear amplitude correction factor is applied to restore true relative echo amplitudes:

$$s_{\text{corr}}(t,k) = 10^{A_k/20} \cdot s_{\text{raw}}(t,k) \tag{3.16}$$

where A_k is the AGC attenuation level in dB for pulse k.

To improve range resolution and facilitate subsequent spectral operations, the raw echoes are resampled by a factor of $\sigma_c = 2$ along the range (fast-time) dimension, increasing the effective sampling frequency from $f_s = 2.8$ MHz to $f_s' = 5.6$ MHz and the number of range samples per echo from 980 to 1960. Resampling is performed via Fourier-domain zero-padding.

The oversampled real-valued echoes are converted to their analytic (complex) representation via the Hilbert transform (Equation **2.6**):

$$\tilde{s}(t,k) = s(t,k) + j \cdot \mathcal{H}\{s(t,k)\} \tag{3.17}$$

This suppresses the negative-frequency spectral image and makes the instantaneous phase and amplitude accessible, which is necessary for the subsequent basebanding and phase correction steps.

Each frequency channel is centered on a carrier frequency $f_0 = 3, 4$ or 5 MHz. To reduce the signal to baseband, the analytic echoes are multiplied by a complex demodulation exponential:

$$s_{\text{bb}}(t,k) = \tilde{s}(t,k) \cdot e^{-j2\pi\Delta f\, t} \tag{3.18}$$

where Δf is the frequency offset between the actual spectral centroid and the desired baseband position. The baseband signal is then bandpass filtered in the frequency domain by applying a rectangular window of width $B = 1$ MHz centered at zero frequency, retaining only the spectral components within the transmitted bandwidth.

As already described in the previous Sections, MARSIS transmits linear frequency-modulated (LFM) chirp pulses of duration $\tau = 250\,\mu$s and bandwidth $B = 1$ MHz. The matched filter reference function (Section **2.2.4**) is the time-reversed conjugated replica of the transmitted chirp. In the frequency domain, the matched filter transfer function is:

$$H_{\text{MF}}(f) = \frac{S^*(f)}{|S(f)|^2 + \epsilon} \tag{3.19}$$

where $S(f)$ is the Fourier transform of the baseband chirp replica, $(\cdot)^*$ denotes complex conjugation, and ϵ is a small regularization constant that prevents division by near-zero values at the band edges. Pulse compression is applied by pointwise multiplication in the frequency domain:

$$S_{\text{pc}}(f,k) = S_{\text{bb}}(f,k) \cdot H_{\text{MF}}(f) \tag{3.20}$$

followed by a Hann spectral window (Equation **2.26**) $W_H(f)$ to suppress range sidelobes, and an inverse Fourier transform to return to the time domain. The resulting range resolution is the one given by Equation **2.42**, which result in 150 m in free space.

Raw MARSIS echoes are not time-aligned along the orbital track: since the spacecraft altitude H_k varies continuously during the observation, the two-way travel time to the surface changes from pulse to pulse, causing the surface return to appear as a diagonal streak in the unprocessed radargram rather than a horizontal line (Figure **3.8a**). To produce a radargram in which the range axis represents true depth below the surface, each echo must be shifted in time by an amount proportional to the deviation of the spacecraft altitude from a common reference value:

$$\delta t_k = \frac{2(H_k - H_{\text{ref}})}{c} \tag{3.21}$$

This shift is applied as a linear phase ramp in the frequency domain, which is equivalent to a sub-sample time translation without interpolation artifacts. After alignment, the surface echo appears as a nearly horizontal feature across the full along-track extent of the radargram (Figure **3.8b**) and the vertical axis can be interpreted directly as two-way travel time measured from a fixed reference altitude ($H_{\text{ref}} = 15$ km above the martian ellipsoid).

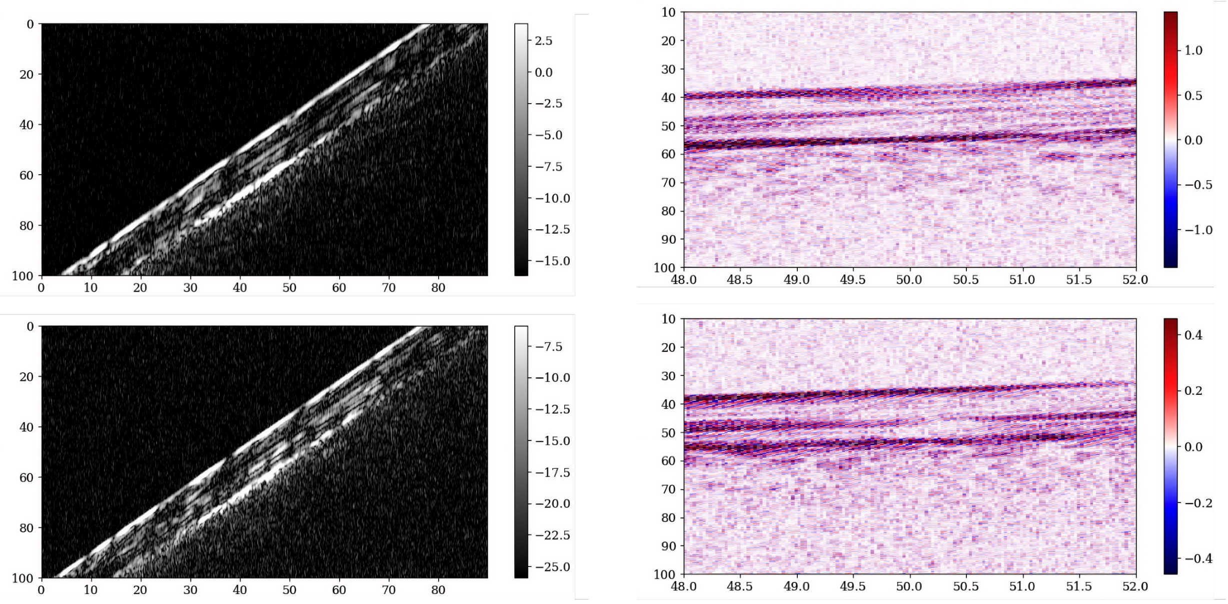

(a) 10711 raw radargrams at 4 and 5 MHz

(b) 10711 radargram at 4 and 5 MHz after topographic alignment and range compression

Figure 3.8: Comparison between 10711 raw and range-compressed radargrams

3.4.2 Generation of simulated data

A coherent electromagnetic simulator of MARSIS surface echoes was developed at INAF to produce synthetic radargrams directly comparable to the processed observations. The simulator takes as input the same orbital and instrumental parameters contained in the L1C data files, that are: spacecraft position and velocity vectors, altitude, sub-spacecraft coordinates, operative frequencies, synthetic aperture sizes, and surface slope estimates, and produces, for each pulse along the track, a simulated range-compressed echo at the same sampling rate and frequency configuration as the real data. The flowchart of the simulator pipeline is shown in Figure **3.9**.

The core of the simulator is a facet-based surface scattering model. For each pulse, the Martian surface topography within a circular region of 100 km radius centered on the sub-spacecraft nadir point is retrieved from the MOLA digital elevation model (MEGDR) at a resolution that adapts with latitude, reaching 512 pixels per degree in the polar regions where the ground track of the observations lies. The retrieved DEM is discretized into a grid of rectangular facets, and for each facet the simulator computes its size, orientation with respect to the spacecraft, and slant range.

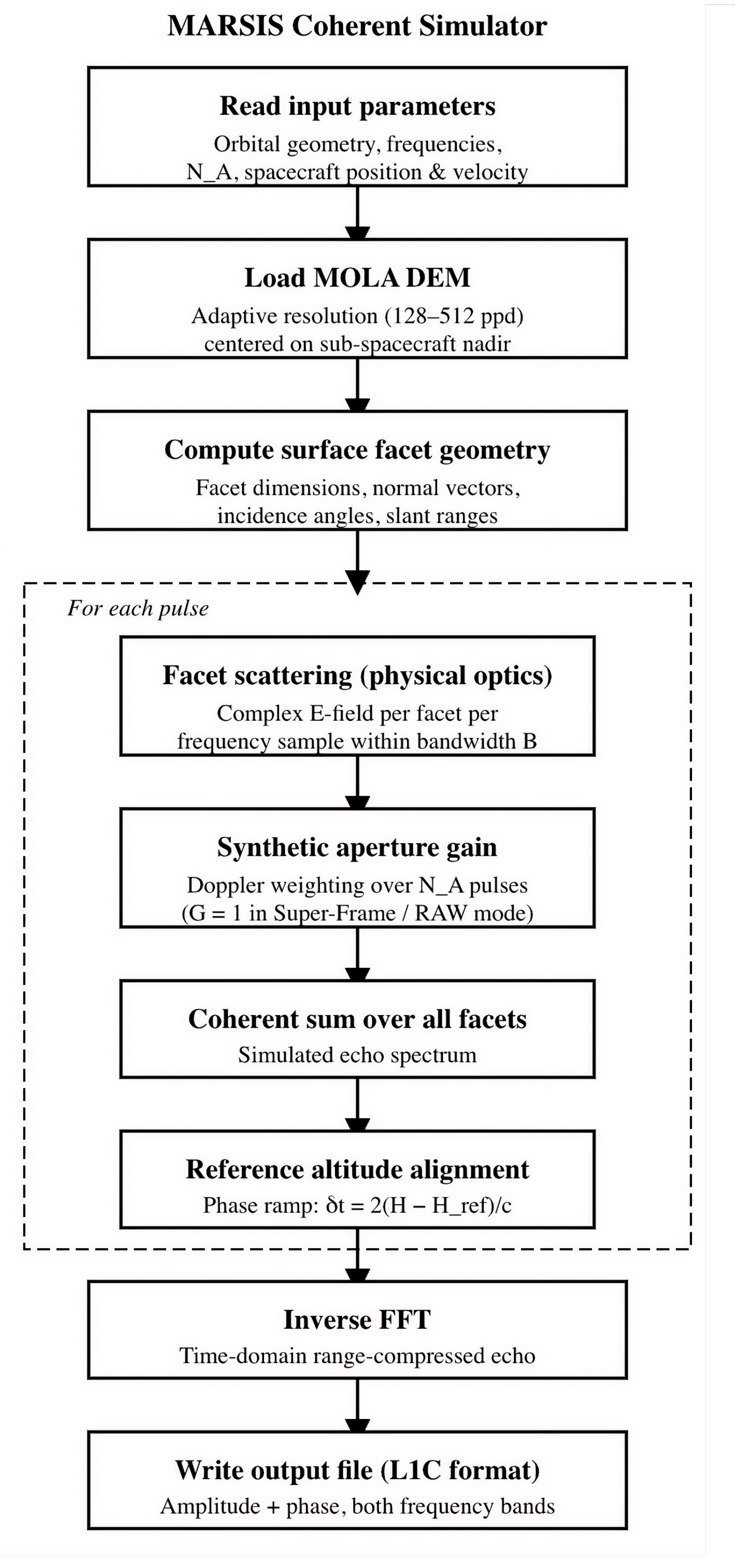

Figure 3.9: Sketch of MARSIS simulator processing pipeline.

The complex electric field scattered back to the spacecraft by each facet is then evaluated using a first-order physical optics model, in which the backscattered amplitude depends on the facet area, the two-way range, and the local incidence geometry through a product of sinc functions of the along- and across-track incidence angles:

$$E \propto \frac{l_a l_b}{R^2} \cos\alpha_x \cos\alpha_y \exp(-2jkR) \cdot \mathrm{sinc}\left(\frac{k}{\pi} l_a \sin\alpha_x\right) \cdot \mathrm{sinc}\left(\frac{k}{\pi} l_b \sin\alpha_y\right) \tag{3.22}$$

where l_a and l_b are the facet half-dimensions, R is the slant range, and α_x, α_y are the local incidence angles in the along- and across-track planes respectively. This expression is evaluated for each frequency sample within the transmitted bandwidth B, and the contributions from all surface facets are coherently summed to form the simulated echo spectrum at that frequency.

The effect of on-board synthetic aperture processing is incorporated through a frequency-domain azimuth gain function computed for each surface facet. This gain reproduces the coherent accumulation of NA pulses performed by MARSIS, accounting for the differential Doppler phase between the actual relative velocity of each surface point and the Doppler centroid assumed during processing, the selected filter offset m and the spacecraft velocity vector. The result is a complex gain pattern that weights the contribution of each facet according to its position relative to the synthetic aperture footprint. In Super-Frame mode, where the raw data are processed on ground without on-board azimuth focusing, the azimuth gain is set to unity and the full unfocused echo is simulated. Finally, a phase ramp is applied in the frequency domain to align all simulated echoes to the same reference altitude H_{ref}, in the same way done in the processing of the real data.

It is important to emphasize what the simulator does not include: it models surface scattering only, with no subsurface returns, no ionospheric propagation effects, no thermal noise, and no instrumental artifacts. This simplicity is precisely what makes the comparison with real data informative. Systematic discrepancies between simulated and observed radargrams, such as timing offsets of the surface echo, data gaps, low-frequency phase modulations, or anomalous echo power at depths inconsistent with surface clutter, cannot be reproduced by the simulator and therefore must originate from physical or instrumental effects absent from the model. As will be shown in the following figures, real MARSIS data exhibit all of these features to various degrees, motivating the correction and characterization steps described in the remainder of this Chapter.

3.5 Challenges in MARSIS data processing

3.5.1 High noise levels in MARSIS data

The five figures shown in Figures **3.10**, **3.11**, **3.12**, **3.13**, and **3.14** illustrate the wide range of data quality encountered accross the MARSIS dataset and motivate the need for careful preprocessing before any quantitative analysis can be performed.

Orbit 10711 (Figure **3.10**) represents the closest approximation to an ideal observation. The acquisition was performed during the Martian night, as indicated by the blue color of the ground track in the Sun elevation angle map, meaning the ionosphere was in its most quiescent state. In both the radargram and the data-vs-simulation overlay, the surface echo in the real data (cyan) is nearly coincident with the simulated surface return (red): the surface appears as a thin, well-defined horizontal feature at approximately 5 km apparent altitude above the ellipsoid, with a morphology that closely matches the prediction from topography alone. The subsurface structure of the SPLD is visible below the surface return, and no gross artifacts are present. This observation represents a benchmark against which the degradation visible in other orbits can be assessed.

Orbit 21571 (Figure **3.11**) is also a high-quality observation in terms of signal-to-noise ratio and echo morphology, but was acquired during daytime (red ground track, positive Sun elevation angle). The effect of the daytime ionosphere is immediately apparent in the overlay panel: the simulated surface echo (red) and the measured surface echo (cyan) are clearly offset in apparent altitude by several kilometers, with the data arriving later than predicted. This systematic delay is the signature of ionospheric dispersion, which retards the group velocity of the radar pulse by an amount proportional to the total electron content along the propagation path. Because the ionosphere is denser during daytime, this delay is significantly larger than in nighttime acquisitions and must be corrected before any depth estimate can be made from the data.

Orbit 02586 (Figure **3.12**) was acquired in standard frame mode rather than Super-Frame mode, as is immediately visible from the periodic vertical striping that interrupts the radargram at regular along-track intervals. These stripes correspond to the boundaries between successive synthetic apertures, where the on-board processing resets and a brief data gap appears. Beyond the frame-boundary artifacts, the comparison with simulations reveals a significantly more complex echo structure than in the nighttime Super-Frame acquisitions: the surface return is broader and less well-defined, off-nadir clutter contributions are visible as diffuse echoes extending well below the surface, and the data-vs-simulation overlay shows substantial mismatch in both position and morphology. The simulation, which assumes perfect synthetic aperture focusing and no ionospheric effect, produces a clean, narrow surface echo that has limited resemblance to the observed radargram.

Orbit 04234 (Figure **3.13**) exemplifies a case of severely degraded data quality. The radargram is dominated by periodic vertical stripes and a high noise floor that partially obscures the surface echo. A faint but spatially coherent echo is nonetheless detectable at approximately ~ 20 km apparent altitude, well below the surface, which is absent from the simulation as expected, since the simulator models surface scattering only. This deep echo is a candidate subsurface return, but its interpretation requires that the strong noise and artifact contamination must be first characterized and suppressed. The overlay panel confirms the poor correspondence between data and simulations, with the surface echo in the data appearing diffuse and displaced relative to the simulated prediction.

Finally, Orbit 04240 (Figure **3.14**) represents the most extreme case of data degradation in this set. The radargram is almost entirely dominated by noise and vertical striping artifacts, to the point where the surface echo is barely distinguishable. The overlay panel shows only the simulated surface return (red) against a featureless noisy background, with virtually no coherent cyan feature identifiable. The ground track map indicates a daytime acquisition at a latitude significantly farther from the pole than the other observations, where ionospheric conditions are typically more severe and variable. This observation illustrates that a non-negligible fraction of MARSIS acquisitions are effectively unusable without aggressive preprocessing, and underscores the importance of quality assessment and selection criteria before any dataset can be used for subsurface characterization.

Taken together, these five examples demonstrate that the gap between ideal simulator predictions and actual MARSIS observations spans a continuum from near-perfect agreement to complete loss of coherent signal. The dominant sources of discrepancy, such as ionospheric delay and dispersion, thermal noise, and frame-mode acquisition artifacts require dedicated non-trivial correction strategies. In any case, as it will be shown in the next sections, there is no perfect correction method that allows to remove completely all the inferences present in MARSIS data.

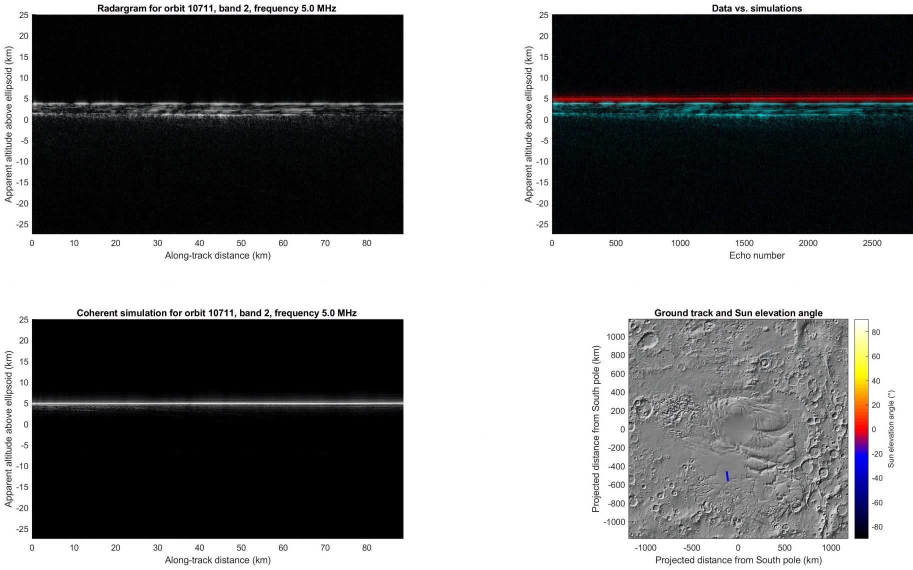

Figure 3.10: Radargram for orbit 10711 at 5 MHz. **Top left**: raw radargram showing the surface echo and subsurface returns as a function of along-track distance. **Top right**: overlay of the measured data (red) and the coherent surface simulation (cyan). **Bottom left**: coherent simulation of the surface return, computed from Mars topography. **Bottom right**: ground track of the orbit (red line) projected onto a map of the Martian south polar region, with the Sun elevation angle indicated by the color scale.

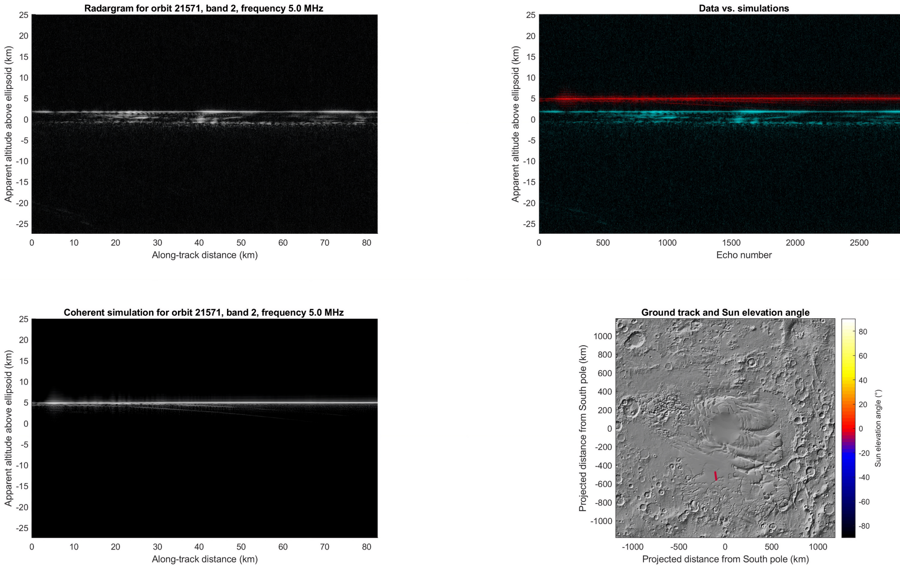

Figure 3.11: Radargram for orbit 21571 at 5 MHz. **Top left**: raw radargram showing the surface echo and subsurface returns as a function of along-track distance. **Top right**: overlay of the measured data (red) and the coherent surface simulation (cyan). **Bottom left**: coherent simulation of the surface return, computed from Mars topography. **Bottom right**: ground track of the orbit (red line) projected onto a map of the Martian south polar region, with the Sun elevation angle indicated by the color scale.

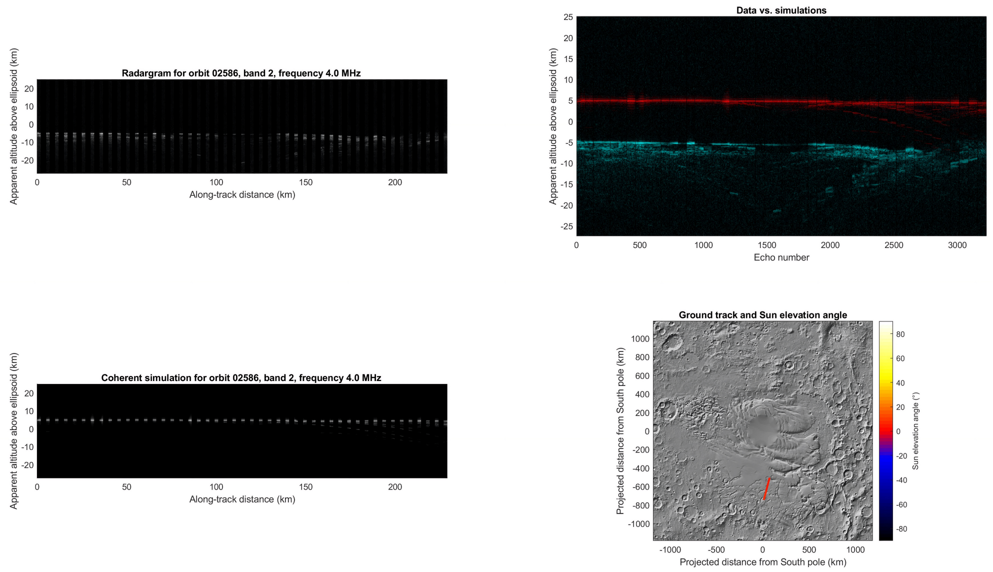

Figure 3.12: Radargram for orbit 02586 at 4 MHz. **Top left**: raw radargram showing the surface echo and subsurface returns as a function of along-track distance. **Top right**: overlay of the measured data (red) and the coherent surface simulation (cyan). **Bottom left**: coherent simulation of the surface return, computed from Mars topography. **Bottom right**: ground track of the orbit (red line) projected onto a map of the Martian south polar region, with the Sun elevation angle indicated by the color scale.

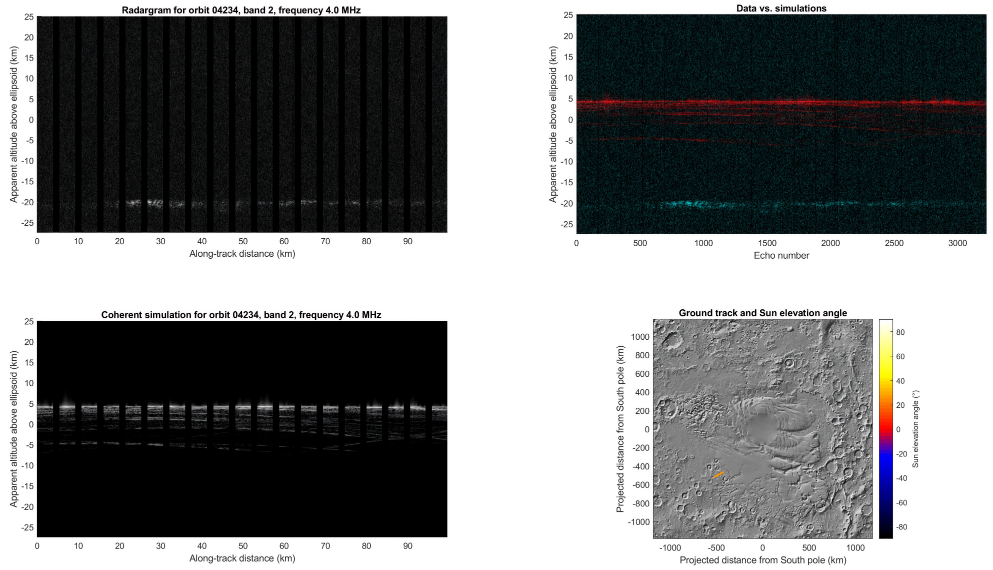

Figure 3.13: Radargram for orbit 04234 at 4 MHz. **Top left**: raw radargram showing the surface echo and subsurface returns as a function of along-track distance. **Top right**: overlay of the measured data (red) and the coherent surface simulation (cyan). **Bottom left**: coherent simulation of the surface return, computed from Mars topography. **Bottom right**: ground track of the orbit (red line) projected onto a map of the Martian south polar region, with the Sun elevation angle indicated by the color scale.

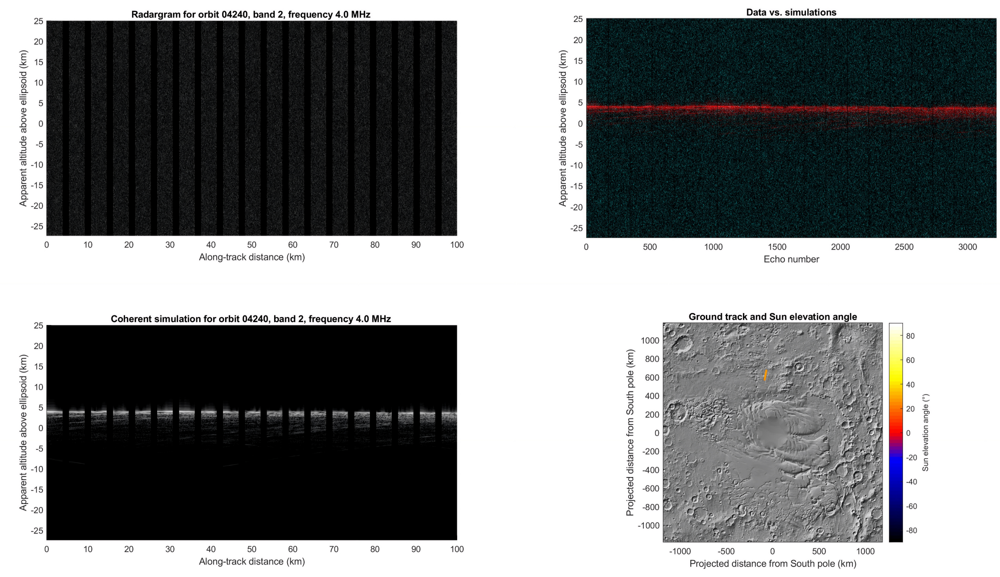

Figure 3.14: Radargram for orbit 04240 at 4 MHz. **Top left**: raw radargram showing the surface echo and subsurface returns as a function of along-track distance. **Top right**: overlay of the measured data (red) and the coherent surface simulation (cyan). **Bottom left**: coherent simulation of the surface return, computed from Mars topography. **Bottom right**: ground track of the orbit (red line) projected onto a map of the Martian south polar region, with the Sun elevation angle indicated by the color scale.

3.5.2 Residual phase delay analysis

The range compression and topographic alignment steps described in Section **3.4.1** produce radargrams in which the surface echo is spatially coherent and the amplitude structure is interpretable. However, a careful inspection of the instantaneous phase of the range-compressed echoes reveals a systematic artifact that must be addressed before any phase-sensitive processing, such as coherent three-dimensional integration, can be performed.

The MARSIS processing pipeline (Section **3.4.1**) demodulates the received signal by mixing it with a local oscillator (Equation **3.18**) and applying low-pass filtering, down-converting the data to an intermediate frequency suitable for digital sampling. This demodulation step introduces a residual phase term that depends on the local oscillator frequency and on the timing of each echo within its pulse repetition interval. The effect is visible in Figure **3.15** (top panel), which shows the instantaneous phase of the range-compressed Super-Frame data for orbit 10737: rather than exhibiting the slowly varying, spatially coherent structure expected from a physical radar echo, the phase oscillates rapidly and apparently chaotically along the echo dimension. This artificial modulation, if uncorrected, would prevent coherent along-track integration and corrupt any phase-based analysis of the data.

An initial correction strategy was developed by simulating the acquisition of the MARSIS transmitted chirp under two conditions: first in true baseband, free of demodulation artifacts, and then in the carrier band, replicating the actual processing chain. Comparing the phase of the surface echo in the two simulated signals allowed the derivation of a filter designed to cancel the artificial phase oscillation. However, a fundamental limitation of this approach is that the derived filter is sensitive to the choice of the model reflector, specifically to its time delay within the recording window. The

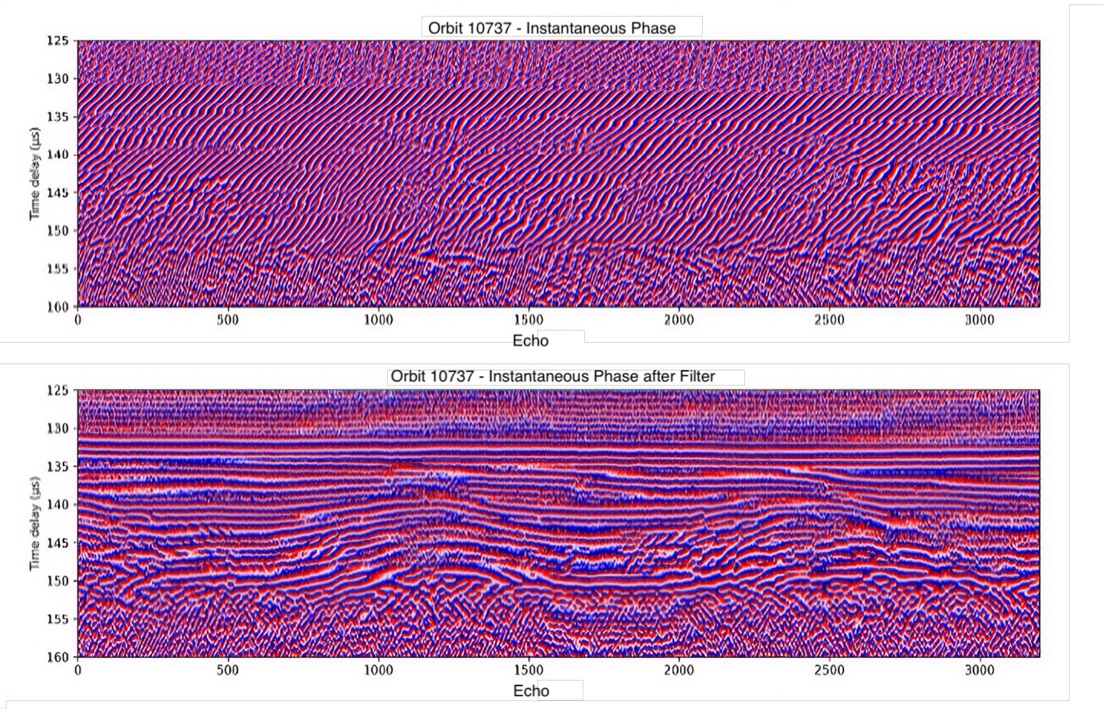

Figure 3.15: Orbit 10737 instantaneous phasegrams. Before filtering (**top panel**), the data phase is dominated by the lateral phase oscillation induced by the "remnant" phase which disappears after filtering (**bottom panel**) to reveal the instantaneous phase of the reflected returns in the data.

effectiveness of the correction therefore depends on the correct selection of this parameter, which cannot be determined unambiguously from the data alone. Furthermore, applying the correction introduced oscillatory behavior in the remodulated waveform that complicated co-registration and caused phase inconsistencies during subsequent processing. For these reasons, this approach was abandoned in favor of a more direct strategy.

The second approach operates directly on the range-compressed data without requiring any model reflector. For each along-track position k, the phase of the surface echo $\phi_s(k)$ is extracted after range compression and converted into an equivalent time delay:

$$\Delta t_\phi(k) = -\frac{\phi_s(k) - \phi_s(0)}{\omega_0} \tag{3.23}$$

being $\omega_0 = 2\pi f_0$ the angular carrier frequency. This measured phase delay is then compared with the two geometric contributions independently known from orbital telemetry: the altitude-dependent delay $\Delta t_H(k)$a and the cumulative radial-velocity delay $\Delta t_r(k)$:

$$\Delta t_H(k) = \frac{2(H_k - H_0)}{c}, \quad \Delta t_r(k) = 2\,\mathrm{PRI}\sum_{i=1}^{k} \frac{v_{r,i}}{c} \tag{3.24}$$

Any residual discrepancy between the measured phase delay and the sum of these two geometric terms must originate from a physical or instrumental effect not accounted for by orbital geometry alone. The result after applying this correction is shown in Figure **3.15** (bottom panel): the instantaneous phase of the surface echo displays a smooth, coherent structure with well-defined horizontal banding reflecting the true phase evolution of the radar wavefront along the orbital track.

Orbit 10737 (Figure **3.16**) was acquired during nighttime at 4 and 5 MHz. The measured phase delays at both frequencies track the spacecraft height delay and radial velocity delay with high fidelity over the full ~ 95 km observation, with residuals of order 1 μs or less. This confirms that for nighttime acquisitions with a quiescent ionosphere, the geometric correction is sufficient and no significant additional delay is introduced by the propagation medium.

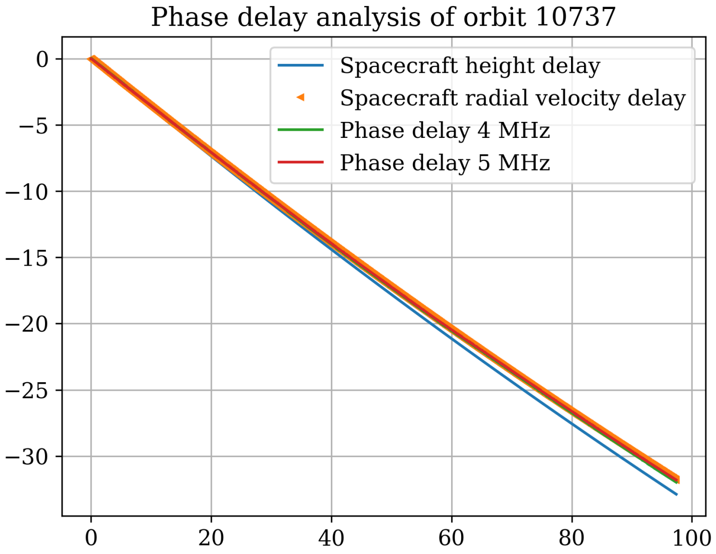

Figure 3.16: Phase delay analysis for orbit 10737. The horizontal axis represents along-track distance in km and the vertical axis the cumulative phase delay in μs relative to the first echo of the observation, with the sign reflecting the direction of spacecraft radial motion along the orbital pass.

Orbit 10711 (Figure **3.17**) presents a more complex picture. At 4 MHz the phase delay follows the geometric prediction closely throughout, while at 5 MHz a progressive divergence develops beyond approximately 60 km along-track. The frequency-dependent nature of the residual suggests a dispersive propagation contribution, though its precise origin requires further investigation.
Orbit 13069 (Figure **3.18**) illustrates the most severe case. From the very beginning of the observation, both frequency bands diverge from the geometric predictions, accumulating residuals of tens of microseconds over the 90 km track. The residual is larger at 4 MHz than at 5 MHz, consistent in principle with the f^{-2} frequency dependence of ionospheric group delay. The irregular step-like features visible around 55-60 km along-track could suggest a strongly structured, rapidly varying ionosphere.
These three examples demonstrate that the phase correction is necessary but not always sufficient. For nighttime acquisitions it fully accounts for the measured surface phase, while for daytime or, more in general, ionospherically disturbed orbits a significant frequency-dependent residual persists.

3.5.3 Ionospheric distortion effects

The Martian ionosphere is a dispersive medium: free electrons produced primarily by solar photoionization slow down and phase-shift electromagnetic waves in a frequency-dependent manner.

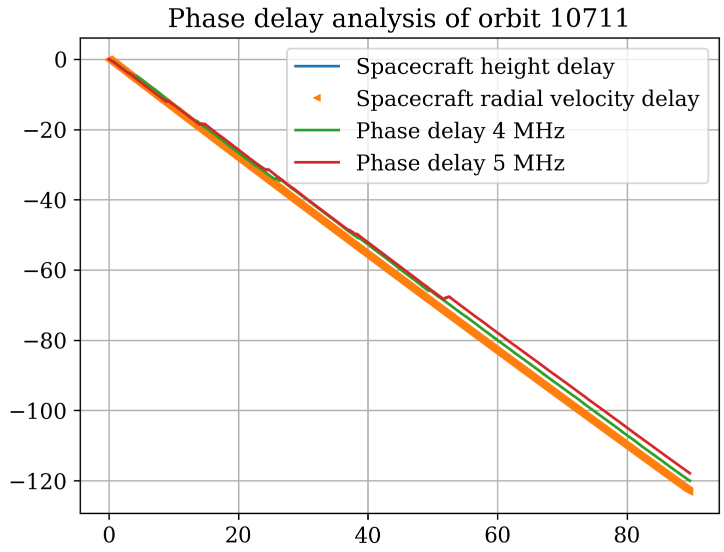

Figure 3.17: Phase delay analysis for orbit 10711.The horizontal axis represents along-track distance in km and the vertical axis the cumulative phase delay in μs relative to the first echo of the observation, with the sign reflecting the direction of spacecraft radial motion along the orbital pass.

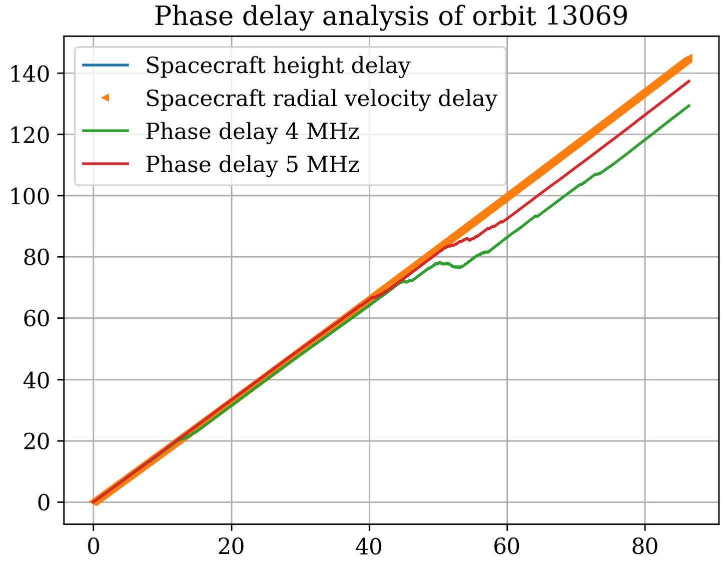

Figure 3.18: Phase delay analysis for orbit 13069. The horizontal axis represents along-track distance in km and the vertical axis the cumulative phase delay in μs relative to the first echo of the observation, with the sign reflecting the direction of spacecraft radial motion along the orbital pass.

When a MARSIS chirp of carrier frequency f_0 and bandwidth $B = 1$ MHz propagates through an ionospheric layer of thickness L, it accumulates a phase shift that can be written as [134, 147]:

$$\Delta\phi(f) = \frac{4\pi f}{c} \int_0^L \left[\sqrt{1 - \frac{f_p^2}{f^2}} - 1 \right] dz \tag{3.25}$$

where $f_p(z) = 8.98\sqrt{N_e(z)}$ Hz is the plasma frequency at altitude z and N_e is the electron density. Because the integrand depends on *f*, different frequency components of the chirp travel at different phase velocities. In the time domain, this dispersive propagation has two principal consequences: the pulse is delayed in its entirety by a group delay (term a_1 in the Taylor expansion of $\Delta\phi$), and its bandwidth is effectively defocused, broadening the range-compressed echo and reducing its peak amplitude (term a_2). Higher-order terms a_3 and a_4 distort the sidelobe structure. The severity of all these effects scales with the total electron content TEC $= \int N_e dz$ and increases sharply as the carrier frequency approaches the peak plasma frequency $f_{p,\max}$.

MARSIS operating frequencies are chosen to remain well above the ionospheric critical frequency precisely to minimize these effects [148]. Nevertheless, even when $f_p \ll f_0$, a residual distortion always remains, and the phase delay analysis described in Section **3.5.2** has already shown that its magnitude varies substantially with solar zenith angle and local ionospheric conditions.

The standard algorithm used to correct ionospheric distortion in MARSIS SS3 data is the Contrast Method (CM), originally developed by [149] and formalized by [147]. The CM operates frame by frame in the frequency domain. Expanding $\Delta\phi(f)$ about f_0 as:

$$\Delta\phi(f) \approx a_0 + a_1(f - f_0) + a_2(f - f_0)^2 + a_3(f - f_0)^3 + a_4(f - f_0)^4 \tag{3.26}$$

the original CM empirically estimates the dominant defocusing coefficient a_2 by iterating over a grid of candidate values and selecting the one that maximizes the total signal energy within the range window. The remaining coefficients a_3 and a_4 are then derived from a_2 through a simplified ionospheric model, while a_1 was addressed in a later version through signal simulation [150].

The correction filter applied to each frame in the frequency domain is:

$$H_{\text{CM}}(f) = \exp\left(-j \cdot \operatorname{sgn}(f_0) \cdot \left[a_2 f^2 + a_3 f^3 + a_4 f^4\right]\right) \tag{3.27}$$

where f is the frequency offset from baseband and $\operatorname{sgn}(f_0)$ takes into account the sign convention at 1.8 MHz versus the higher frequencies. The loop iterates over 200 candidate a_2 values on the first frame of each orbit, then refines to 40 iterations around the previous estimate for subsequent frames, exploiting along-track continuity of the ionosphere.

Despite nearly two decades of development, complete correction of ionospheric distortion in MARSIS data remains an open problem. The most recent study on the subject [148] introduced a New CM that replaces the empirical a_2 estimation with ionospheric models, one for daytime and nighttime each. The New CM slightly improves performance, particularly near the terminator and at the lowest operative frequencies, and provides more accurate TEC estimates on the nightside. However, a critical finding of that work is that achieving good ionospheric correction, as measured by range compression energy, does not uniquely constrain the TEC: different models with different physical assumptions can produce indistinguishable range compression quality while yielding substantially different TEC values. On the dayside, where the condition $f_p \ll f_0$ is less well satisfied, the various existing algorithms diverge considerably and no consensus has been reached.

For the present work, the CM implementation applied to the Ultimi Scopuli dataset is the version described by [147, 150] and implemented in the INAF pipeline. The phase delay analysis of Section **2.37** already provides a direct empirical assessment of the residual ionospheric contribution orbit by orbit: nighttime acquisitions show residuals consistent with zero excess delay after geometric correction, confirming that the CM is effective under those conditions, while daytime orbits retain a measurable frequency-dependent residual that the CM cannot fully remove. This motivates restricting the primary subsurface analysis to nighttime data, for which the ionospheric contribution is demonstrably small.

3.6 Noise Reduction Algorithm

The range-compressed radargrams (L1C data) produced at the end of the pipeline described in Section **3.4.1** are complex matrix of dimensions $N_t \times N_x$, where $N_t = 980$ is the number of range samples and N_x is the number of along-track echoes. In this domain, thermal noise and other incoherent contributions are distributed nearly uniformly across all range samples and all along-track positions. The coherent surface and subsurface echoes, by contrast, are spatially structured: they are continuous along track and occupy a well-defined band in both range frequency and Doppler frequency. Exploiting this structure in the two-dimensional Fourier domain provides a natural basis for noise suppression.

The full procedure is described in [151]. A 2D discrete Fourier transform is applied to the complex echo matrix, mapping it from the range sample, along-track echo domain to the range frequency, Doppler frequency domain. The range frequency axis reflects the spectral content of the range-compressed chirp, which for the processed L1C data, already baseband-shifted during ground processing, is confined between $[-B/2, B/2]$. The Doppler frequency axis reflects the along-track phase coherence of the echoes: a spatially continuous return concentrates energy into a narrow set of Doppler columns near zero Hz, while incoherent noise populates the full Doppler bandwidth uniformly.

A first spectral masking step retains only range frequencies within $[-0.5, 0.5]$ MHz, discarding out-of-band components without affecting the chirp signal.

Figure **3.19** shows the result of this operation for orbit 10737 at 4MHz. The upper-left panel shows the 2D spectrum before filtering: the signal occupies a horizontally extended band at mid-range frequencies, while energy is distributed across the full Doppler extent. The median amplitude per Doppler column (upper-right panel) reveals a sharply peaked structure centered near 4 MHz, corresponding to the coherent surface echo, sitting above a nearly flat noise floor that spans the remaining ~ 120 Hz of Doppler bandiwdth. This contrast between signal and noise in the Doppler domain is the foundation of the filtering step. For each Doppler frequency column j of the 2D spectrum, the median spectral amplitude across all range frequency rows is computed. Columns whose median amplitude falls below a data-adaptive threshold are set to zero; only columns carrying significant coherent energy are retained:

$$m_j = \text{median}_i \, |\tilde{P_{ij}}| \tag{3.28}$$

$$\tilde{P_{ij}} = \begin{cases} P_{ij} & \text{if } m_j \geq \theta \\ 0 & \text{otherwise} \end{cases} \tag{3.29}$$

The threshold θ is determined automatically via Otsu's method [152] applied to the normalized median amplitude distribution. The method treats the distribution of m_j values as a one-dimensional

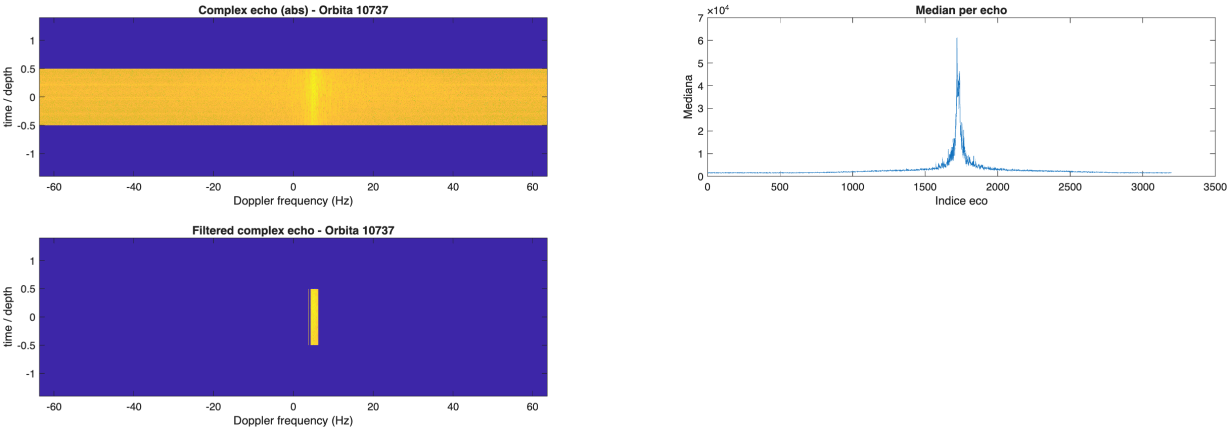

Figure 3.19: **Top panel, left**: 2D spectrum of radargram 10737 at 4 MHz. **Top panel, right**: Median amplitude per Doppler column. *Bottom panel*: Filtered 2D spectrum.

histogram and finds the threshold that maximizes the between-class variance of the two resulting groups, which are columns classified as signal and columns classified as noise. The method requires no manual tuning and adapts to the SNR conditions of each orbit independently. The lower panel-left panel of Figure **3.19** hows the 2D spectrum after filtering: only the narrow coherent peak near zero Doppler is retained, and the broad incoherent floor is completely suppressed.
The filtered 2D spectrum $\tilde{P}$ is then transformed back to the range sample, along-track echo domain via the inverse 2D FFT, yielding the noise-reduced complex echo matrix. Figure **3.20** shows the radargrams of orbit 10737 before (left) and after (right) filtering. Before filtering, a uniform noise floor fills the entire range window above and below the surface echo, with the subsurface SPLD layering only marginally visible against the background. After filtering, the noise floor above the surface echo is substantially reduced, the surface return sharpens, and the faint subsurface structure below sample ~ 450 becomes more clearly resolved. To quantify the improvement introduced by

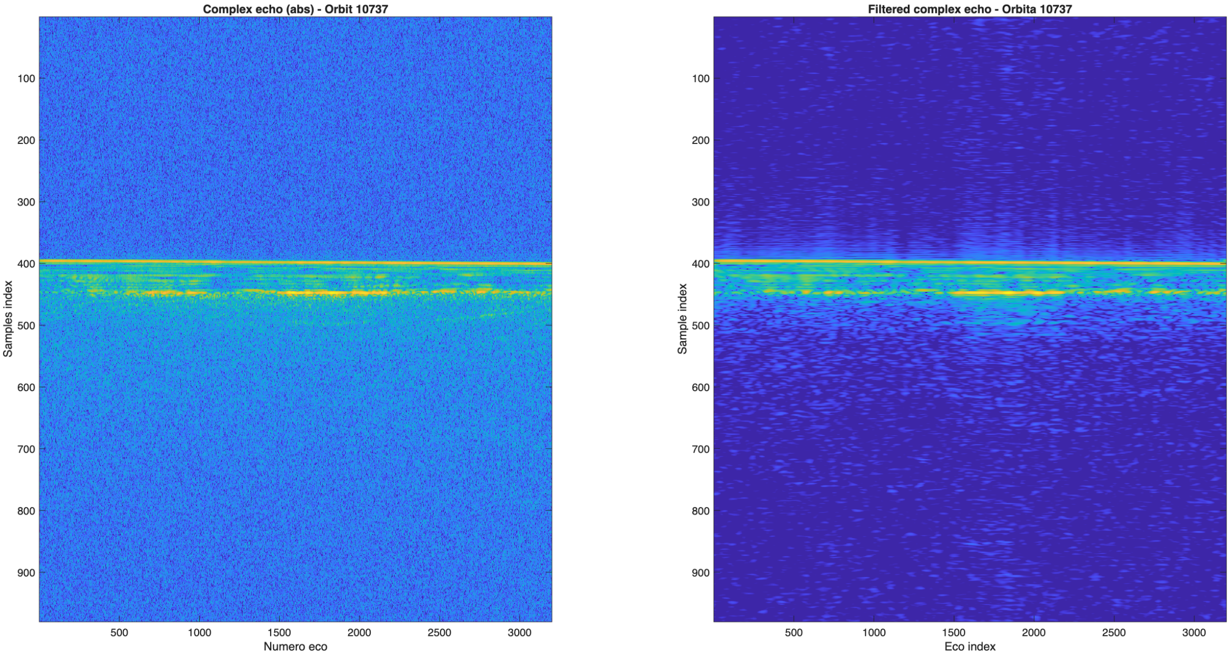

Figure 3.20: **Left**: radargram of orbit 10373. **Right**: Filtered radargram of orbit 10737

the filter, an SNR estimate is computed independently for each along-track echo before and after filtering. For a given column k of the complex echo matrix, the signal power is estimated as the mean power of the top $K = 5$ range samples ranked by amplitude:

$$P_{\text{sig}}(k) = \frac{1}{K}\sum_{i=1}^{K} |s_{(i),k}|^2 \tag{3.30}$$

The noise floor is estimated as the median power of all samples whose power falls at or below the 50th percentile of the column:

$$P_{\text{noise}}(k) = \text{median}\ \left\{|s_{i,k}|^2 : |s_{i,k}|^2 \leq P_{50}(k)\right\} \tag{3.31}$$

The per-column SNR(k) = $\log[P_{\text{sig}}(k)/P_{\text{noise}}(k)]$, and a scalar summary is obtained by averaging over all valid columns. This estimator requires no assumption about the position of the surface echo within the range window and adapts automatically to the local noise level of each echo.
Table **3.6** reports the SNR estimates before and after filtering for all orbits processed in this work. The filter consistently improves the SNR across the entire dataset, with a mean gain of approximately 16.8 dB regardless of carrier frequency or acquisition conditions. The improvement is remarkably uniform, ranging from ~ 12 dB for the most degraded orbits (21578, 15136) to ~ 21 dB or the cleanest ones, confirming that the Otsu-based Doppler thresholding adapts effectively to varying noise conditions. The post-filtering SNR falls uniformly in the $33 - 48$ dB range across all orbits and both frequency bands, providing a consistent and well-characterized input for the subsequent subsurface analysis.

3.7 2D focusing results and comparison

In Chapter **2** we presented several focusing algorithms, from SAR to migration, that differ in the operator they implement (closed-form inverse $\mathcal{A}^{-1}$, or adjoint $\mathcal{A}^{*}$) and in the assumptions required for their validity. Before describing the implementation adopted in this work, it is necessary to justify why the adjoint-based backprojection is the only approach that can be applied to MARSIS data in a scientifically defensible way. The answer follows directly from the data quality analysis carried out in the preceding Sections.
All migration algorithms operates under the BA (Section **2.5.2**). Both Stolt migration and omega-k (Section **2.4.3**) implement a closed-form of the inverse operator $\mathcal{A}^{-1}$ by exploiting the fact that, under specific geometric assumptions, the forward operator becomes diagonal in the Fourier domain. We recall that the diagonalization is met in case of: a flat acquisition surface at constant altitude; a spatially uniform and exactly known propagation velocity; and zero-offset geometry. None of these conditions is met by MARSIS. The spacecraft altitude varies continuously along the orbit due to the elliptical orbit of Mars Express, introducing position-dependent phase errors that destroy the Fourier-domain diagonality. More critically, the propagation path is not uniform: the signal travels through the Martian ionosphere, whose electron density varies in space and time and whose dispersive effect depends on the total electron content in a way that cannot be reduced to a single equivalent velocity. As demonstrated in Section **3.5.3**, residual phase delays of several microseconds persist even after the CM correction, with frequency-dependent and orbit-dependent behavior. For Stolt and omega-k, such phase errors do not merely degrade the image: they accumulate coherently across the full aperture in the wavenumber domain, producing large-scale artifacts that can completely mask the true subsurface signal. The assumption of a known, smooth velocity model required by omega-k algorithm is equally unfeasible: the effective propagation speed varies

Table 3.6: SNR before and after 2D Fourier domain noise filtering for all processed orbits. Values are averaged over all along-track echoes. ΔSNR denotes the improvement introduced by the filter.

Orbit	f_0 **(MHz)**	$\mathbf{SNR_{pre}}$ **(dB)**	$\mathbf{SNR_{post}}$ **(dB)**	**ΔSNR (dB)**	**Orbit**	f_0 **(MHz)**	$\mathbf{SNR_{pre}}$ **(dB)**	$\mathbf{SNR_{post}}$ **(dB)**	**ΔSNR (dB)**
10711	4	27.3	43.1	15.8	14967	3	28.8	44.0	15.3
10711	5	25.4	40.6	15.3	14967	4	30.2	46.6	16.5
10737	4	30.7	46.1	15.3	15055	3	27.6	43.7	16.1
10737	5	28.0	43.2	15.2	15055	4	28.6	46.5	17.9
10961	4	21.9	40.6	18.7	15110	3	23.0	39.3	16.3
10961	5	19.7	38.3	18.6	15110	4	27.7	44.1	16.4
12685	4	27.9	46.4	18.5	15136	3	26.2	38.1	11.9
12685	5	25.2	43.2	18.0	15136	4	27.4	40.3	12.8
12759	4	20.4	39.1	18.7	15198	3	22.5	38.9	16.4
12759	5	20.7	38.3	17.6	15198	4	28.8	46.3	17.5
12766	4	25.4	44.0	18.5	17021	3	23.1	39.2	16.2
12766	5	24.7	42.5	17.8	17021	4	26.8	44.1	17.2
12814	4	26.8	44.0	17.2	17116	4	31.5	47.2	15.8
12814	5	25.8	41.7	16.0	17116	5	29.3	44.7	15.4
12840	4	31.2	47.6	16.4	17299	4	29.3	47.8	18.4
12840	5	29.2	44.2	14.9	17299	5	26.6	43.8	17.2
12895	4	26.3	45.5	19.1	19216	4	24.8	44.0	19.3
12895	5	25.3	43.3	18.0	19216	5	26.4	43.8	17.4
12969	4	30.0	47.1	17.1	19283	4	27.7	45.0	17.3
12969	5	28.7	43.9	15.3	19283	5	28.3	44.2	15.9
13043	4	19.7	37.2	17.6	19350	4	24.4	40.3	16.0
13043	5	21.4	38.1	16.8	19350	5	25.8	42.1	16.3
13050	4	23.2	42.0	18.8	19357	4	24.7	44.4	19.6
13050	5	24.0	41.4	17.4	19357	5	25.9	43.3	17.4
13069	4	24.4	39.9	15.6	19491	4	20.6	37.1	16.6
13069	5	23.0	38.5	15.5	19491	5	21.3	38.6	17.2
14853	4	26.0	43.7	17.7	19498	4	21.1	42.0	20.9
14853	5	23.9	41.7	17.8	19498	5	21.7	41.6	19.9
14879	4	28.3	44.2	15.9	21571	4	19.7	37.9	18.3
14879	5	27.2	43.0	15.7	21571	5	19.0	36.4	17.5
					21578	4	18.6	33.1	14.5
					21578	5	18.2	33.0	14.8

Mean: $\text{SNR}_{\text{pre}} = 25.4\,\text{dB}$ $\text{SNR}_{\text{post}} = 42.2\,\text{dB}$ $\Delta\text{SNR} = 16.8\,\text{dB}$

with ionospheric TEC along the orbit in a way that cannot be characterized with sufficient accuracy from the available data.

Phase-shift migration (Section **2.4.3**) implements the adjoint $\mathcal{A}^*$ through iterative backward wave propagation in the frequency domain, and is therefore conceptually closer to backprojection. However, it still requires a layered velocity model $c(z)$ and a flat acquisition surface, and it assumes one-way propagation. The layered velocity assumption fails for the same reason as above: the ionosphere introduces a horizontally variable, frequency-dependent phase contribution that cannot be collapsed into a depth-only function $c(z)$. Furthermore, phase-shift migration is sensitive to the accuracy of the velocity model at each depth level: errors in the assumed $c(z)$ cause the wavefield to be back-propagated along incorrect trajectories, producing focusing errors that grow with depth. In our problem, where the subsurface targets of interest are located at ~ 1.5 km depth beneath the surface, even small velocity errors in the SPLD layer ($\epsilon_s \pm \delta\epsilon$) would translate into significative range positioning errors.

The adjoint operator $\mathcal{A}^*$ implemented as time-domain or frequency-domain (see in the following) backprojection (Sections **2.4.2** and **2.6.2**) operates only under the BA condition. Thus, it requires no knowledge of the propagation medium beyond the two-way travel time from each radar position to each image point. This makes it uniquely robust to the data quality issues identified in the preceding sections. Residual ionospheric phase errors, provided they are small enough not to decorrelate the signal across the aperture, result in a mild defocusing of the PSF rather than in coherent artifacts, because each aperture position contributes independently to the sum and errors do not accumulate in a structured way. Also, the algorithm is inherently robust to non-uniform along-track sampling, which arises naturally from the variable PRF of MARSIS: since the contribution of each echo is evaluated at its exact geometric position, irregularly spaced apertures do not introduce grating lobes or aliasing.

The validity of BA implies that, as discussed in Section **2.6.2**, the backprojection provides an adjoint solution rather than a true inverse: it does not recover quantitatively accurate reflectivity amplitudes, and its PSF is the autocorrelation of the system response rather than a delta function. For the purpose of this work, that is, detecting and mapping the spatial extent of the bright basal reflector and selecting which orbits carry sufficient coherent signal for 3D reconstruction, and considering the quality of MARSIS data, this limitation is acceptable.

Finally, the 2D focusing step described here is implemented via backprojection rather than conventional SAR focusing algorithms (Section **2.3**) deliberately, for reasons of consistency with the 3D reconstruction pipeline. Since the multi-orbit 3D reconstruction necessarily uses backprojection, as argued above, applying the same operator in the 2D preparatory stage ensures that the SNR gain and phase behavior observed on individual orbits are directly comparable to what will be obtained in the full 3D combination, with no change in algorithm, assumptions, or image-domain geometry between the two stages. Hence, it should be noted that conventional SAR focusing would equally represent a valid and well-established choice for the 2D case alone, and is indeed applied to MARSIS data both on-board and in ground processing pipelines; the preference for backprojection here is driven purely by pipeline consistency rather than by any fundamental inadequacy of SAR-based approaches for single-orbit focusing.

3.7.1 2D frequency-domain backprojection

The backprojection integral can be evaluated either in the time domain or in the frequency domain, with practical differences. In the time-domain formulation, the coherent summation is performed by interpolation, which is computationally expensive and introduces resampling errors that degrade phase fidelity. Also, it must be applied independently at each depth, making the algorithm scale poorly with image size.

The frequency-domain formulation avoids interpolation entirely. Since the Fourier transform of a time-shifted signal $s(t - \tau_0)$ is $S(f)e^{-j2\pi f \tau_0}$, the phase compensation for any desired delay $\tau(x_i, z_j)$ can be applied exactly as a complex exponential multiplication in the frequency domain, without any resampling. This is both more accurate, since the phase is preserved at every frequency, and more efficient, because the FFT of each echo is computed once and reused for all image depths. The approach implemented here follows the Sum and Delay with Shifting and Zoom method in the frequency domain ([153, 154], Figure **3.21**).

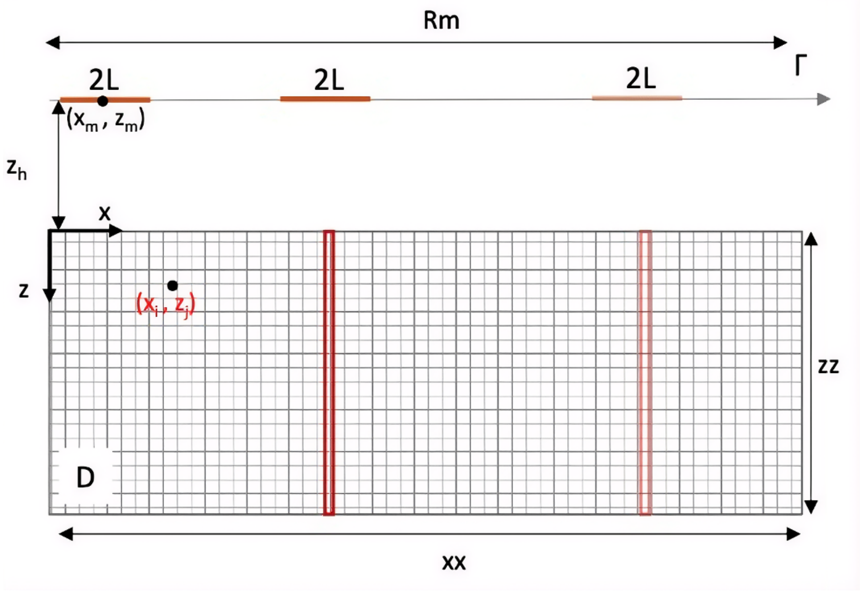

Figure 3.21: Geometrical sketch of the Shift and Zoom algorithm. The investigation domain D has dimensions xx and zz. The acquisition domain is Rm, with 2L being the zoom window dimension.

Starting from the complex L1C range-compressed echo matrix of dimensions 980 × 3200, the algorithm proceeds as follows. The along-track distance of each radar position from the first echo is computed from the three-dimensional position vector:

$$x_m = \sqrt{(X_m - X_0)^2 + (Y_m - Y_0)^2} \tag{3.32}$$

using the spacecraft position vector encoded in the L1C file. A 1D FFT is applied along the range dimension to obtain the frequency-domain data $S(f, x_m)$, and a bandpass selection retains only the frequency samples within the chirp bandwidth, corresponding to indices *if1=315* to *if2=666* out of 980 samples at $f_s = 2.8$ MHz, yielding an absolute frequency vector $f_1 = f_{\text{baseband}} + f_0$ spanning approximately ±0.5 MHz around the carrier f_0 (Figure **3.22**). For each image pixel (x_i, z_j) the

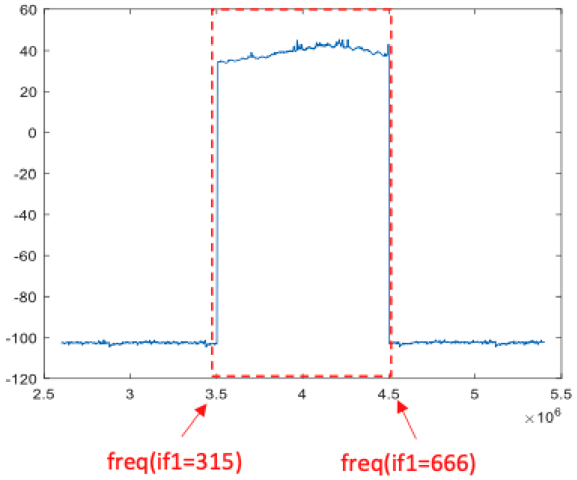

Figure 3.22: Selection of the frequency domain for the shift and zoom algorithm

focused reflectivity is computed by coherent summation over a sliding aperture of half-length $L = 1200$ pulses centered at the output position:

$$\rho(x_i, z_j) = \sum_{if=1}^{N_f} \sum_{m=1}^{N_m} \text{DATAF}(if, (x_m, z_m)) \cdot \exp\left\{ j \cdot 2 \cdot k_{eq}(if, z_j, z_h) \cdot \sqrt{(x_i - x_m)^2 + (z_j - z_m)^2} \right\} \tag{3.33}$$

where k_eq is the equivalent wavenumber that accounts for the mixed propagation path: free-space propagation from the spacecraft to the surface over altitude z_h, and propagation through the SPLD ice layer from the surface to the image point at depth z_j:

$$k_{eq}(if, z_j, z_h) = k_0(if)\sqrt{\epsilon_{eq}}, \quad \epsilon_{eq} = \left(\frac{\sqrt{\epsilon_s} z_j + z_h}{z_j + z_h} \right)^2 \tag{3.34}$$

where $k_0(if)$ is the free-space wavenumber at frequency f, and $\epsilon_s = 3.1$ is the relative permittivity of the SPLD. The equivalent permittivity ϵ_{eq} is a depth-weighted harmonic mean of two media: it reduces to unity for $z_j = 0$ (pure free-space propagation). In the limit $z_j \ll z_h$, which is representative of MARSIS observations, ϵ_{eq} remains close to unity with a small correction proportional to z_j/z_h, reflecting the fact that most of the propagation path occurs through free space. In the opposite limit $z_j \gg z_h$, ϵ_{eq} tends to ϵ_s, correctly recovering full in-ice propagation.

The half-aperture parameter L controls the number of echoes contributing to each image pixel and therefore directly controls both the along-track resolution and the coherent integration gain. Figure **3.23** shows the effect of varying L on the focused image of orbit 10737 at 4MHz, for $L = 5, 50, 500$ and 1200 echoes respectively.

At $L = 5$, the aperture is too short to produce any meaningful focusing: the difference between the unfocused and the focused image is negligible. At $L = 50$ some bright spots at range distance ~ 23 start to become visible, but numerous scattering artifacts remain. At $L = 500$ the focusing is substantially improved: the surface echo is sharp and the deeper reflectors are clearly visible. The transition from $L = 500$ to $L = 1200$ yields a further improvement in image cleanliness and subsurface contrast, confirming that the data remain phase-coherent over apertures of at least 1200 echoes on each side of the output point, but the qualitative gain is modest compared to the step from $L = 50$ to $L = 500$.

The choice of aperture length involves a trade-off between coherent integration gain and output coverage. With $N_x = 3200$ echoes per orbit, the focusing runs from echo L to echo $N_x - L$, producing a valid focused output only over central $N_x - 2L$ echoes. With $L = 1200$ this yields only ~ 800 echoes, which may be insufficient to cover the full extent of the target region when the reflector lies near the edge of the ground track. With $L = 500$ the valid window grows to ~ 2200 echoes, comfortably encompassing a region of ~ 50 km, while still providing a substantial integration gain as demonstrated above.

In terms of physical scale, assuming the typical MARSIS pulse repetition frequency and orbital velocity, the half-apertures $L = 500$ and $L = 1200$ correspond to along-track integration distances of approximately 14 km and 33 km respectively. Since the backprojection algorithm explicitly compensates the geometric phase at each pixel, integration lengths exceeding the Fresnel radius of $\sim 7 - 10$ km (see Eq. **1.59**) do not introduce incoherent summation, provided phase coherence is maintained across the aperture. The theoretical azimuth resolutions achieved are approximately 800 m ($L = 500$) and 340 m ($L = 1200$), representing an improvement of roughly one order of magnitude over the $7 - 10$ km footprint of standard unfocused SAR processing. The modest improvement from $L = 500$ to $L = 1200$ suggests that residual phase errors, primarily of ionospheric origin, begin to limit the effective coherence beyond $\sim$14 km, preventing the algorithm from fully exploiting the larger aperture. For the purpose of orbit selection, where spatial coverage of the target is as important as focusing quality, $L = 500$ therefore represents the more appropriate choice. The purpose of this 2D focusing step in the processing pipeline is to provide a quantitative criterion

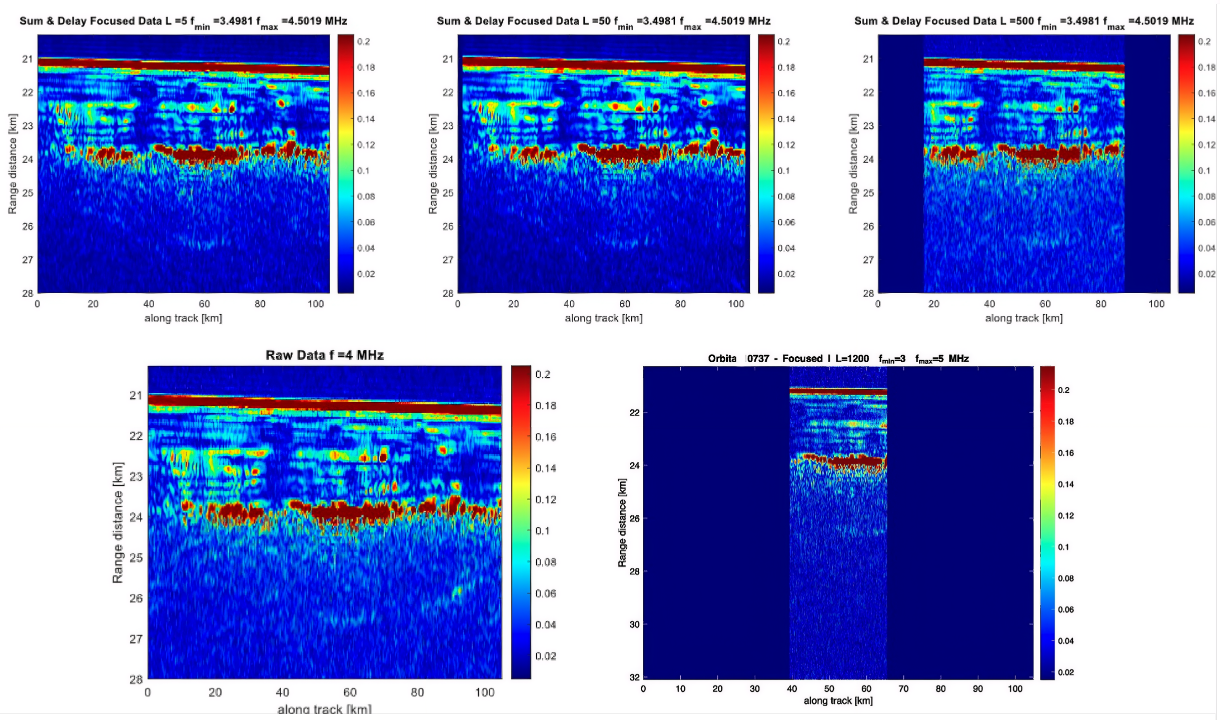

Figure 3.23: Comparison between different choices of the half-aperture parameter L for the radargram of orbit 10737 at 4 MHz: L=5 (top-left), L=50 (top-middle), L=500 (top-right) and L=1200 (bottom-right).

for orbit selection prior to the full 3D backprojection. An orbit whose data are coherent, that is, whose surface and subsurface echoes have stable phase along track and have been adequately corrected for ionospheric delay and residual phase errors, will show a measurable SNR improvement after focusing: incoherent noise does not focus, while coherent echoes from genuine reflectors do.

Conversely, an orbit with severe ionospheric distortion, strong frame-mode artifacts, or insufficient SNR will show little or no improvement after focusing. The SNR metric defined in Section **3.6** is computed for both the unfocused and focused radargrams, and the gain ΔSNR serves as the selection threshold: only orbits exceeding a minimum gain are retained for the multi-orbit 3D reconstruction described in the following.

3.7.2 Orbit selection for 3D processing

The starting point for orbit selection is the identification of all available MARSIS orbits whose ground track passes within a geographically meaningful distance of the bright basal reflector at Ultimi Scopuli. From the full MARSIS archive, 31 orbits were identified whose ground track falls within approximately 50 km of the reflector location (see Table **3.6**). Among them, a subset of 10 orbits passes in proximity of the bright reflector, providing the direct nadir coverage necessary for subsurface sounding at the target depth.

The top of Figure **3.24** shows the ground tracks of all 31 orbits (red lines) overlaid on the MOLA topography of the study area, displayed in orthographic projection centered on the South Pole. The coordinate axes are expressed in kilometers in the projected plane. The bottom of Figure **3.24** displays the cumulative Fresnel zone footprint of the full orbit set, with the white region indicating surface area sampled by at least one orbit within one Fresnel zone radius, while the top of Figure **3.24** overlays the same ground tracks on the MOLA shaded relief, providing geographic context. The dominant orientation of the ground tracks is northeast–southwest, nearly parallel to one another, reflecting the near-polar orbital geometry of Mars Express. This quasi-parallel configuration has direct consequences for the angular diversity of the dataset and hence for the tomographic reconstruction quality (see Section **2.7.2**). The Fresnel zone radius for each orbit position is computed as:

$$r_F = R_{\text{Mars}} \cdot \arccos\left(\frac{r_{sc}^2 + R_{\text{Mars}}^2 - (h + \lambda/4)^2}{2R_{\text{Mars}} r_{sc}}\right) \tag{3.35}$$

where r_{sc} is the spacecraft distance from the Martian center and h= $r_{sc} - R_{\text{Mars}}$ is the altitude above the surface. At typical MARSIS altitude of 300 km, $f_0 = 4$ MHz, it yields to $r_F \approx 9$ km, defining the minimum cross-track spacing between orbits below which their Fresnel zones begin to overlap and provide redundant surface coverage.

The 2D frequency-domain backprojection algorithm was applied to the original L1C range-compressed data first, prior to any other phase correction or noise filtering for all ten orbits (see Table **3.7**) identified close to the bright basal reflector (83° S, 193° E). The reason is that the reconstruction algorithm and the MARSIS data processing pipelines were developed in parallel: the phase correction and noise filtering procedures were implemented at a later stage, after the initial 3D reconstruction analysis had already begun. As a consequence, the SNR gain after focusing at this stage reflects the intrinsic phase coherence of each orbit on uncorrected data, and provides a direct diagnostic of data quality independent of any preprocessing. The algorithm was run with half-aperture $L = 500$ echoes, yielding a valid focused output window of approximately 68 km per orbit, that is enough to encompass the full extent of the reflector region while still providing a substantial coherent integration gain. The SNR metric defined in Section **3.6** was computed on the raw data (all echoes) and on the focused data (valid columns only, i.e. columns $L + 1$ to $N_x - L$), and the gain ΔSNR= $\text{SNR}_{\text{focused}} - \text{SNR}_{\text{raw}}$ was used as the selection criterion. The results are summarized in Table **3.7**. Of the ten candidate orbits, only three show a positive ΔSNR: orbits

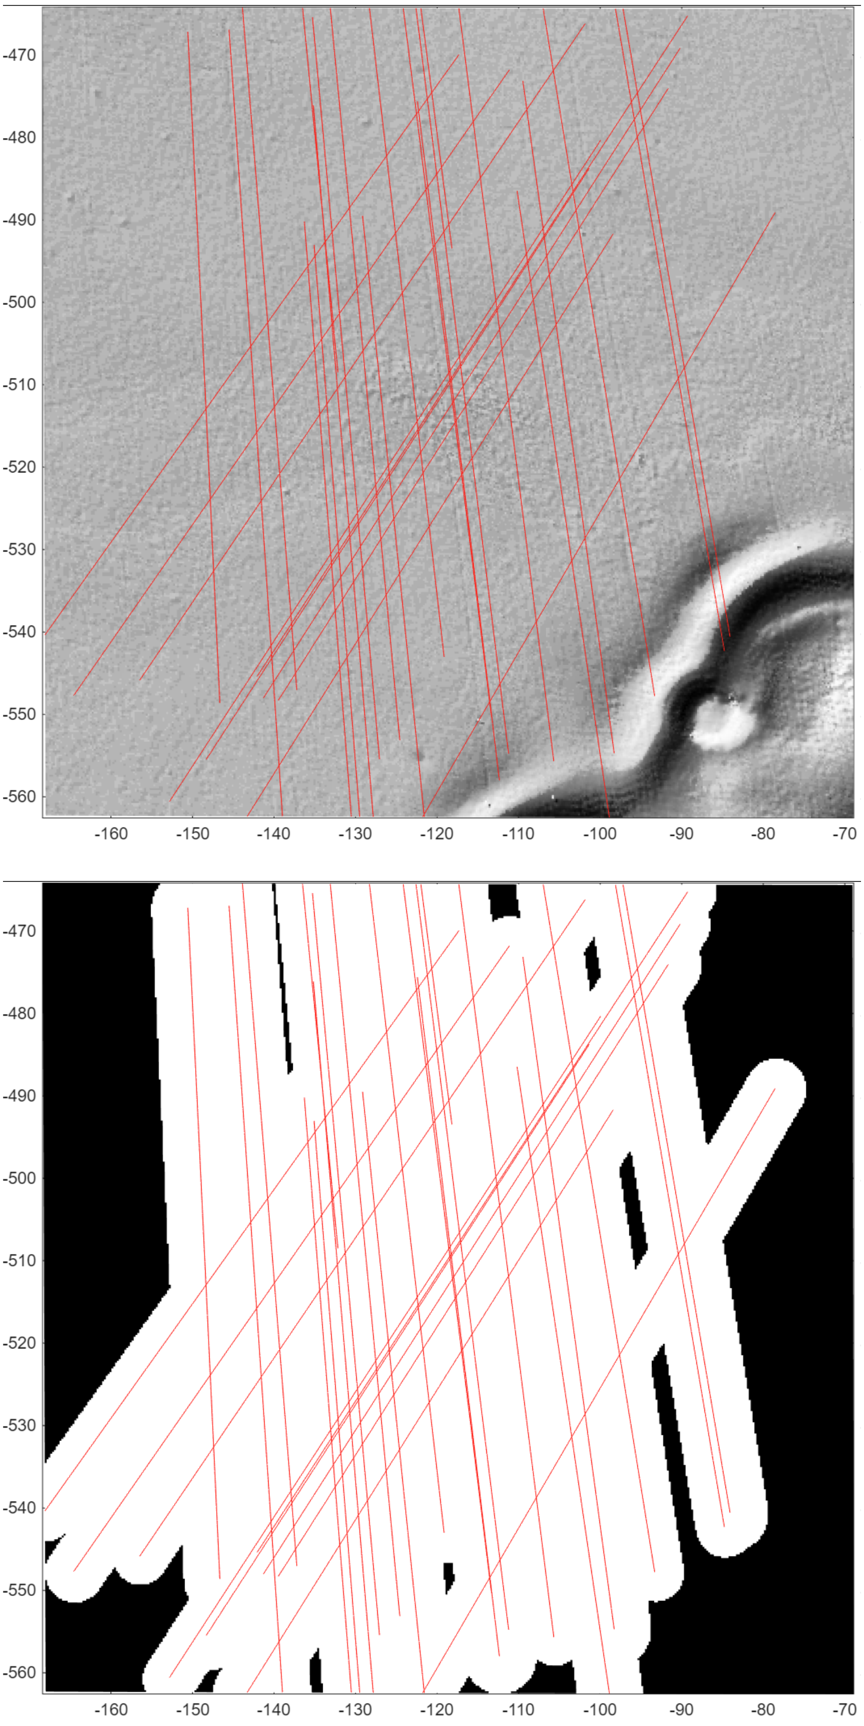

Figure 3.24: **Top**: Ground tracks of all orbits in a 50 km radius around the study area. **Bottom**: Cumulative Fresnel zone footprint of the full orbit set.

10737, 12840, 15110 (Figures **3.23** and **3.25**). These are precisely the orbits whose ground tracks pass closest to the bright reflector. The remaining seven orbits all show a degradation in SNR after focusing, indicating that the data are not sufficiently phase-coherent across the $L = 500$ aperture to produce a net coherent integration gain, likely due to a combination of stronger ionospheric phase errors, less favorable acquisition geometry, or lower raw signal level at those ground track positions. These seven orbits are therefore excluded from the 3D reconstruction. The three selected orbits form a geometrically compact but angularly limited dataset. The ground tracks of two of them are approximately parallel. As discussed in Section **2.7.2**, this falls in the regime of partial quasi-tomographic reconstruction: sufficient to partially constrain the 3D structure of the reflector, but leaving significant gaps in k-space that will result in anisotropic resolution. The implications for the 3D image quality are discussed in Section **3.8**.

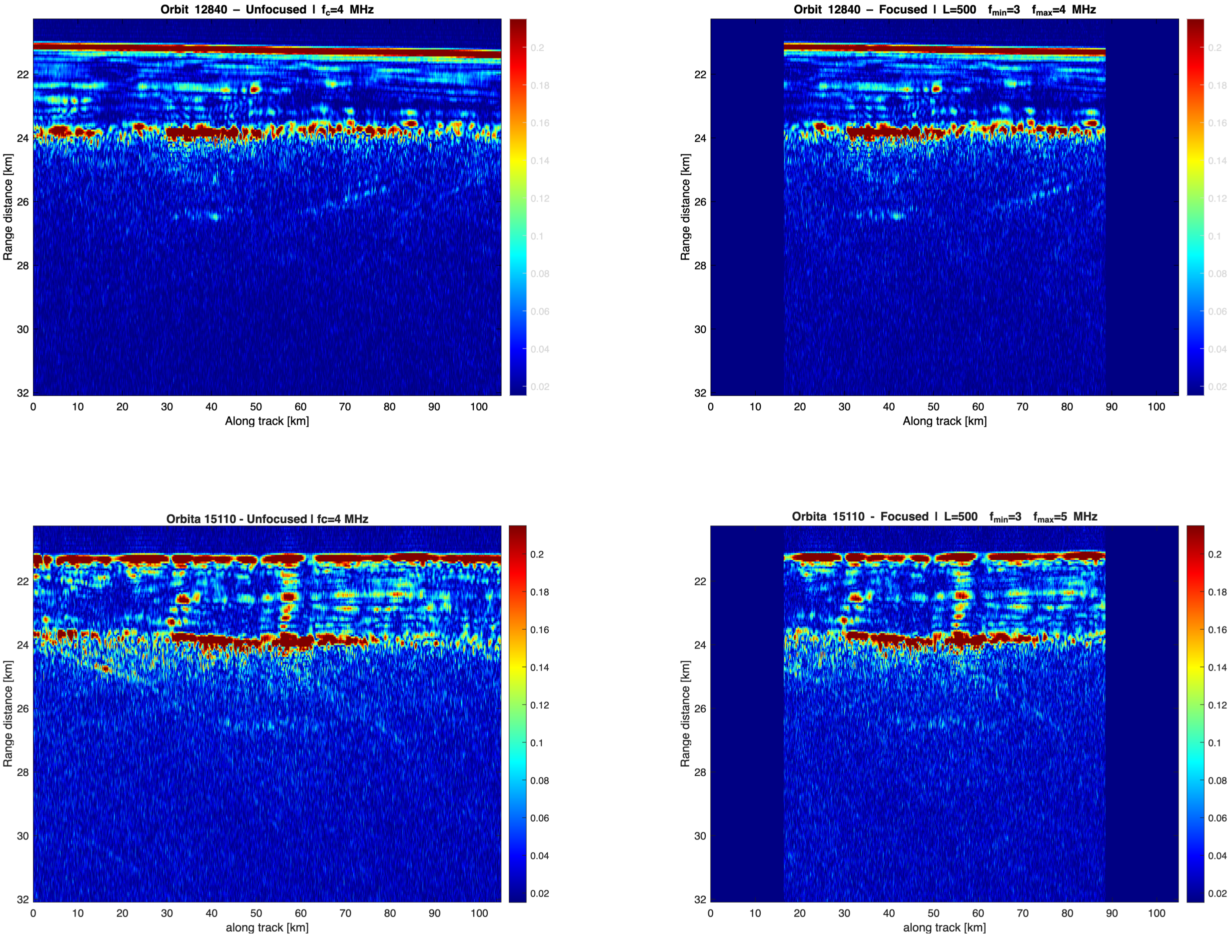

Figure 3.25: **Top**: Unfocused vs focused radargram of orbit 15110 at 4 MHz. **Bottom**: Unfocused vs focused radargram of orbit 15110 at 4 MHz.

Following the noise filtering algorithm described in Section **3.6**, the focusing algorithm was applied anew to the full set of 31 orbits passing over the lake candidate region, rather than the three-orbit subset considered in the preliminary analysis. For each orbit, the noise-corrected complex echoes at 4 MHz were used as input, with the SNR values at this stage corresponding to the post-filtering values reported in Table **3.6**. The extended orbit set and the improved input data quality are expected to yield a more coherent three-dimensional reconstruction, with better suppression of sidelobe artifacts and improved sensitivity to subsurface reflectors. Table **3.8** summarizes the orbits

Table 3.7: SNR before and after 2D frequency-domain backprojection focusing for the 10 orbits passing within 50 km of the bright basal reflector at Ultimi Scopuli. $\mathrm{SNR}_{\mathrm{pre}}$ is taken from Table **3.6**. $\mathrm{SNR}_{\mathrm{focused}}$ is computed on the valid aperture window (columns $L+1$ to $N_x - L$). Orbits enclosed between thick rulers are retained for 3D reconstruction.

Orbit	f_0 **(MHz)**	$\mathbf{SNR}_{\mathbf{pre}}$ **(dB)**	$\mathbf{SNR}_{\mathbf{focused}}$ **(dB)**	**ΔSNR (dB)**
10711	4	27.3	25.4	-1.9
10711	5	25.4	24.1	-1.3
10737	4	30.7	32.1	1.4
10737	5	28.0	28.6	0.6
10961	4	21.9	20.8	-1.1
10961	5	19.7	17.9	-1.8
12685	4	27.9	26.4	-1.5
12685	5	25.2	23.1	-3.1
12840	4	31.2	33.0	1.8
12840	5	29.2	31.3	2.1
13069	4	24.4	24.0	-0.4
13069	5	23.0	22.1	-0.9
14853	4	26.0	24.7	-1.3
14853	5	23.9	22.4	-1.3
14879	4	28.3	28.1	-0.2
14879	5	27.2	26.5	-0.7
15110	3	23.0	23.2	0.2
15110	4	27.7	28.1	0.4
15136	3	26.2	25.4	-0.8
15136	4	27.4	26.6	-0.8

used in this second focusing run together with their input SNR. The difference with respect the original dataset is evident: among the 31 orbits, only 4 show a negative ΔSNR, leaving 27 orbits usable for the extended 3D reconstruction.

Table 3.8: SNR before and after 2D frequency-domain backprojection focusing for all 31 orbits processed with noise-corrected data. $\mathrm{SNR}_{\mathrm{pre}}$ corresponds to $\mathrm{SNR}_{\mathrm{post}}$ from Table **3.6**, i.e. the SNR of the noise-filtered echoes provided as input to the focusing algorithm. $\mathrm{SNR}_{\mathrm{focused}}$ is computed on the valid aperture window (columns $L+1$ to $N_x - L$). All values refer to the 4 MHz band. Orbits marked with † were discarded from the 3D reconstruction due to insufficient focusing gain (ΔSNR).

Orbit	**$\mathrm{SNR}_{\mathrm{pre}}$ (dB)**	**$\mathrm{SNR}_{\mathrm{focused}}$ (dB)**	**ΔSNR (dB)**
10711	43.1	45.0	1.9
10737	46.1	47.0	0.9
10961	40.6	41.2	0.6
12685	46.4	47.4	1.0
12759	39.1	40.2	1.1
12766	44.0	44.5	0.5
12814	44.0	44.8	0.8
12840	47.6	48.1	0.5
12895	45.5	46.8	1.3
12969	47.1	48.0	0.9
13043	37.2	38.5	1.3
13050	42.0	42.4	0.4
13069	39.9	40.5	0.6
14853	43.7	45.1	1.4
14879	44.2	45.1	0.9
14967	46.6	48.2	1.6
15055	46.5	47.2	0.7
15110	44.1	46.1	2.0
15136	40.3	40.7	0.4
15198	46.3	47.0	0.7
17021	44.1	44.9	0.8
$17116^{\dagger}$	47.2	45.6	-1.6
17299	47.8	49.0	1.2
$19216^{\dagger}$	44.0	42.8	-1.2
19283	45.0	45.7	0.7
19350	40.3	41.4	1.1
19357	44.4	45.2	0.8
19491	37.1	37.3	0.2
19498	42.0	43.3	1.3
$21571^{\dagger}$	37.9	37.0	-0.9
$21578^{\dagger}$	33.1	31.4	-1.7

3.8 3D Reconstruction Results

3.8.1 3D frequency-domain backprojection

The 2D focusing algorithm described in Section **3.7.1** operates on a single orbit at a time, integrating echoes along the along-track direction only. The extension to three dimensions requires combining data from multiple orbits with different ground tracks, so that the coherent summation spans both the along-track and cross-track directions simultaneously. As established in Section **2.7.2**, this multi-orbit configuration is what provides the angular diversity necessary for tomographic reconstruction of the subsurface volume.

The 3D algorithm retains the same Sum and Delay with Shifting and Zoom structure in the 2D scenario [153, 154], but the image domain is now a three-dimensional volume and the measurement points span all echoes from all orbits simultaneously. The three orbits selected in Section **3.7.2** are loaded jointly and their frequency-domain data stacked into a single measurement matrix of dimension $N_f \times (N_x \cdot N_{\text{orbit}})$, with $N_x = 3200$ and $N_{\text{orbit}} = 3$.

The reconstruction volume is defined as a Cartesian grid centered on the bright basal reflector, with spatial steps $\Delta x = \Delta y = 30$ m in the horizontal plane and depth samples running from range index $it_1 = 392$ to $it_2 = 592$, corresponding to a depth window of approximately 5 km below the surface. The horizontal extent of the grid is:

$$x \in [-580, -545]\,\text{km}, \quad y \in [-140, -120]\,\text{km} \tag{3.36}$$

in the ortographic projection centered on the South Pole, yielding a domain of approximately 35×20 km around the reflector location.

For each image pixel (x_i, y_j, z_k), the algorithm selects all measurement points within an integration radius $R_{\text{int}} = 1200$ m of the horizontal projection of the pixel. It is worth noting an important conceptual difference between the 2D and 3D implementations. In the 2D algorithm of Section **3.7.2**, the aperture parameter L was defined as a number of echoes, a purely along-track quantity controlling how many consecutive pulses contribute to each image pixel. In this 3D algorithm, the analogous parameter R_{int} is instead a physical radius in the horizontal plane: all measurement points from any orbit whose ground projection falls within a circle of radius R_{int} around the image pixel (x_i, y_j) contribute to the coherent sum, regardless of which orbit they belong to or their along-track position. This formulation is the natural generalization to the multi-orbit case, where there is no single along-track axis and contributions from different orbits must be weighted by geometry rather than by index. The value $R_{\text{int}} = 1200$ m represents a trade-off between coherent integration quality and computational cost: larger integration radii increase the number of contributing measurement points quadratically, making the 3D reconstruction over a domain of several kilometers computationally prohibitive. The chosen value ensures a sufficient number of contributing points per pixel while keeping the processing time tractable.

The algorithm computes the coherent sum:

$$\rho(x_i, y_j, z_k) = \sum_{m \in \mathcal{N}(x_i, y_j)} \omega_m \cdot S(f, \mathbf{r}_m) \cdot \exp\left\{ j \cdot 2 \cdot \frac{2\pi f}{c} \cdot R_{3D} \right\} \tag{3.37}$$

where $S(f, \mathbf{r}_m)$ is the complex amplitude spectrum of the range-compressed echo at measurement point $\mathbf{r}_m = (x_m, y_m, z_m)$, with $z_m = -z_h$ being the spacecraft altitude, and $R_{3D} = \sqrt{(x_i - x_m)^2 + (y_j - y_m)^2 + (z_k - z_m)^2}$ is the three-dimensional range from measurement point $\mathbf{r}_m$ to image pixel. The outer sum runs over all frequencies $f \in [f_{\min}, f_{\max}]$ and the inner sum runs over all measurement points m within the integration radius. The reconstruction is performed for each pixel (x_i, y_j, z_k), with z_k spanning the full subsurface depth range of interest, yielding a three-dimensional volume $\rho(x_i, y_j, z_k)$. Note that z_k represents the subsurface image depth, distinct from the spacecraft altitude $z_m = -z_h$.

The term ω_m is an angle-dependent amplitude weight:

$$\omega_m = \cos^2 \theta_m, \quad \cos \theta_m = \frac{|z_k - z_m|}{R_{3D} + \epsilon} \tag{3.38}$$

where θ_m is the off-nadir angle between the vertical and the line connecting the measurement point to the image pixel. This two-way amplitude weighting accounts for the antenna gain pattern of MARSIS, strengthening near-nadir returns and attenuating contributions from large off-nadir angles where both the antenna gain and the signal-to-clutter ratio are reduced. The reconstruction was tested with several values of the exponent: $\alpha = 1$, corresponding to the standard Kirchhoff obliquity factor; $\alpha = 2$, corresponding to the two-way amplitude antenna gain; and $\alpha = 4$, tested as an empirical value providing stronger off-nadir suppression. No significant differences were observed in the reconstructed volumes across these cases. This insensitivity is attributable to two factors: the integration radius $R_{\text{int}} = 1200$ m limits contributions from large off-nadir angles, and the target region is topographically flat, making oblique reflections negligible in practice. Since $\alpha = 2$ carries a clear physical interpretation as the two-way amplitude antenna gain, it is adopted as the final value.

Prior to 3D focusing, the data of each orbit are normalized to their respective maximum amplitude, so that no single orbit dominates the coherent sum due to differences in absolute signal level between acquisitions. The frequency band used for focusing corresponds to indices of the baseband spectrum, selecting a bandwidth of approximately 1 MHz around the 4 MHz carrier.

3.8.2 3D results

Figure **3.27** shows a horizontal slice of the reconstructed volume at depth $z = 2.944$ km. This depth reported here is an apparent depth, computed assuming that the signal propagates at the speed of light in vacuum. The true depth is related to the apparent one by $d_{\text{true}} = d_{\text{app}}/\sqrt{\epsilon_r}$, where ϵ_r is the permittivity of the overlying medium. Assuming $\epsilon_r \approx 3.5$, consistent with Martian ice and regolith (with a dust mixture of about $15 - 20\%$), the corrected depth is approximately 1.5 km, in agreement with the results reported in [143]. Figure **3.26** shows three horizontal slices at apparent depths $z \sim 0$, ~ 3 and ~ 5.5 km, providing a qualitative view of the signal distribution across the reconstructed volume.

The most evident feature in Figure **3.26** is that the backprojected amplitude is not uniformly distributed across the investigation domain, but is strongly concentrated along the two ground tracks of the contributing orbits. This is consistent with the geometry of the problem, since for a given pixel, the coherent sum in Equation **3.37** only receives contributions from measurement points within the integration radius $R_{\text{int}} = 1200$ m of the pixel's horizontal projection. Pixels far from any orbit track receive no contribution and remain at zero amplitude, while pixels close to a ground track accumulate the maximum number of coherent contributions. The result is that the reconstructed image, at this stage, retains a strong imprint of the acquisition geometry.

The signal amplitude reaches its maximum in the vicinity of the orbit intersection, at approximately $(x, y) \approx (-558, -129)$ km, corresponding to the bright reflector location. Away from the intersection, the image collapses to a one-dimensional structure aligned with each orbit track, reflecting the fact that a single orbit provides focusing only in the range and along-track directions, leaving the cross-track dimension unresolved.

A rough estimate of the along-track extent of the reconstructed signal can be obtained from this picture. The bright region along the more energetic orbit track extends over approximately $20 - 25$ km in the along-track direction, consistent with the values estimated by [143], while the cross-track width of the illuminated strip is approximately $2 - 3$ km. The stacked view in Figure **3.26** confirms that the signal is depth-localized: the two shallowest slices show structured

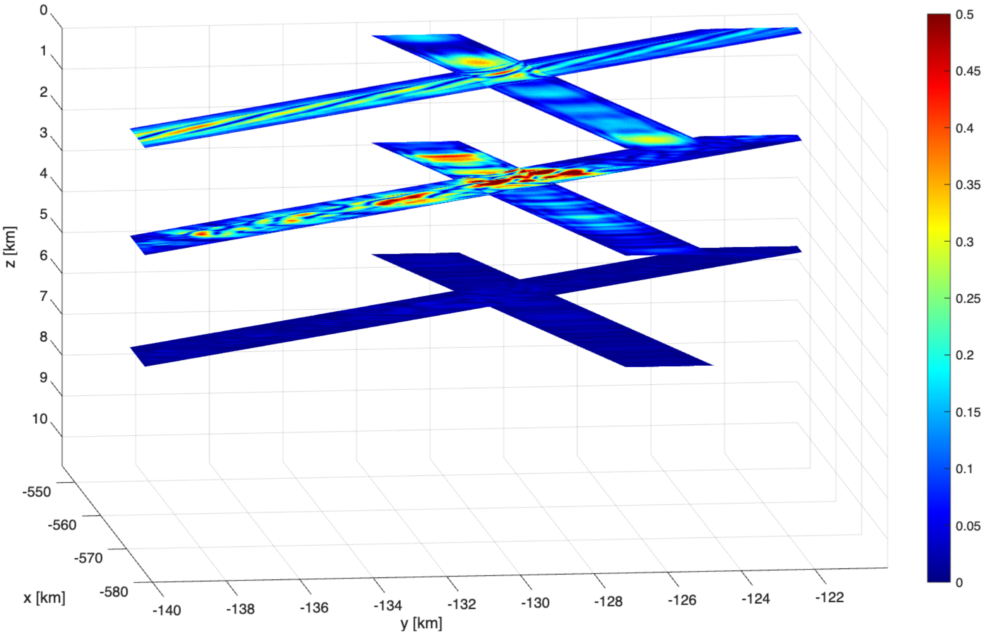

Figure 3.26: Slices of the 3-orbits reconstruction volume at ~ 0, ~ 3 and ~ 5.5 km depth. From L1C data.

amplitude along the orbit tracks, while the deepest slice is essentially featureless, consistent with the absence of a strong reflector at that depth. The intermediate slice at $z \approx 2.9$ km shows the strongest response, in agreement with the expected apparent reflector depth. Several limitations of the present reconstruction must be noted. First, these results are obtained from L1C data, that are range-compressed echoes without ionospheric phase correction or noise filtering. The residual phase errors of Section **3.5.2** are expected to partially decorrelate the coherent sum across orbits, reducing the effective focusing gain and broadening the point spread function. Second, the three orbits used here share a predominantly NE–SW orientation so the angular diversity in the cross-track direction is limited. A true tomographic reconstruction would require orbits spanning a

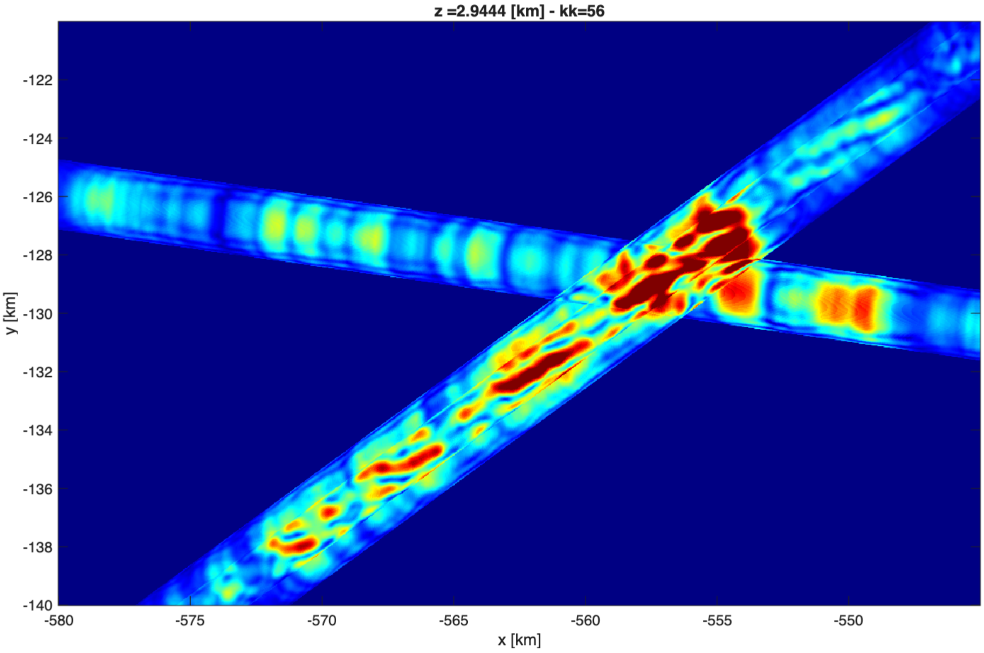

Figure 3.27: Slice of the 3-orbits reconstruction volume at ~ 3 km depth. From L1C data.

wider range of azimuth angles, which is not available in the present dataset. Also, we should specify again that the backprojection result therefore represents a blurred and geometrically distorted estimate of the true subsurface reflectivity, rather than a quantitatively accurate reconstruction. In other words, the backprojection algorithm provides a qualitative image of the subsurface reflectivity distribution, but not a quantitative one. As a consequence, no absolute reflectivity values are extracted from the 3D reconstruction, and comparisons between pixels at different locations or depths should be interpreted with caution.

Next, we applied the same backprojection algorithm to the noise and phase corrected dataset. We generated the L2C data by applying the ionospheric correction (Section **3.5.3**), then we applied the phase correction (Section **3.5.2**) and the noise reduction algorithm (Section **3.6**).

A comparison of the horizontal slices at the apparent reflector depth ($z \approx 2.94$ km, Figures **3.27** and **3.28**) reveals a qualitatively consistent picture in terms of overall geometry: in both cases, the backprojected amplitude is concentrated along the two orbit ground tracks, with a maximum near their intersection at $(x, y) \approx (-558, -129)$. However, several differences are apparent. In the L1C reconstruction, the amplitude distribution along the main orbit track is strongly fragmented, with multiple disconnected high-amplitude patches of comparable intensity distributed over the full extent of the track. In the phase-corrected reconstruction, this fragmentation is substantially reduced: the bright region consolidates into a more continuous and coherent strip, with a well-defined amplitude maximum near the orbit crossing point. This behavior is consistent with the expected effect of ionospheric phase correction: residual phase ramps across the aperture introduce destructive interference that breaks the coherent sum into isolated constructive patches, whereas after correction the phase history is better aligned across orbits, and the coherent integration produces a spatially more compact response.

A further notable feature of the corrected reconstruction is the appearance of a spatially isolated and well-focused bright spot ai $x \sim 550$, $y \sim 130$ km. This feature is present but poorly defined in the L1C result, where it is partially masked by the higher and more uniform noise floor. Its emergence after phase and noise correction suggests that it corresponds to a genuine coherent reflector whose phase coherence is partially destroyed by the uncorrected ionospheric delay and the higher noise level in the L1C data. The stacked depth slices (Figures **3.26** and **3.29**) confirm these observations at the volumetric level. The phase and noise corrected slices show a lower background level, a sharper depth localization of the signal at the intermediate apparent depth ($z \approx 3 - 5$ km), and reduced along-track spreading of the response. The shallowest and deepest slices remain essentially featureless in both cases, consistent with the absence of strong reflectors at those depths. Quantitatively, the maximum normalized amplitude in the corrected reconstruction is comparable to the L1C case, as expected given that both datasets are normalized to their respective maxima prior to backprojection. The key improvement is therefore not in peak amplitude but in image contrast and spatial coherence: the corrected reconstruction concentrates the available signal energy into fewer, more localized pixels, which is the signature of improved focusing. These results confirm that phase and noise correction, while not altering the fundamental geometry of the reconstruction, produces a measurable improvement in image quality. The reduction of incoherent phase errors across orbits increases the effective coherent integration gain, reduces the noise floor of the reconstructed volume, and improves the spatial localization of subsurface reflectors. The limitations discussed before, i.e. limited angular diversity, anisotropic resolution, absence of absolute reflectivity calibration, remain valid for both reconstructions and are inherent to the acquisition geometry rather than to the processing chain.

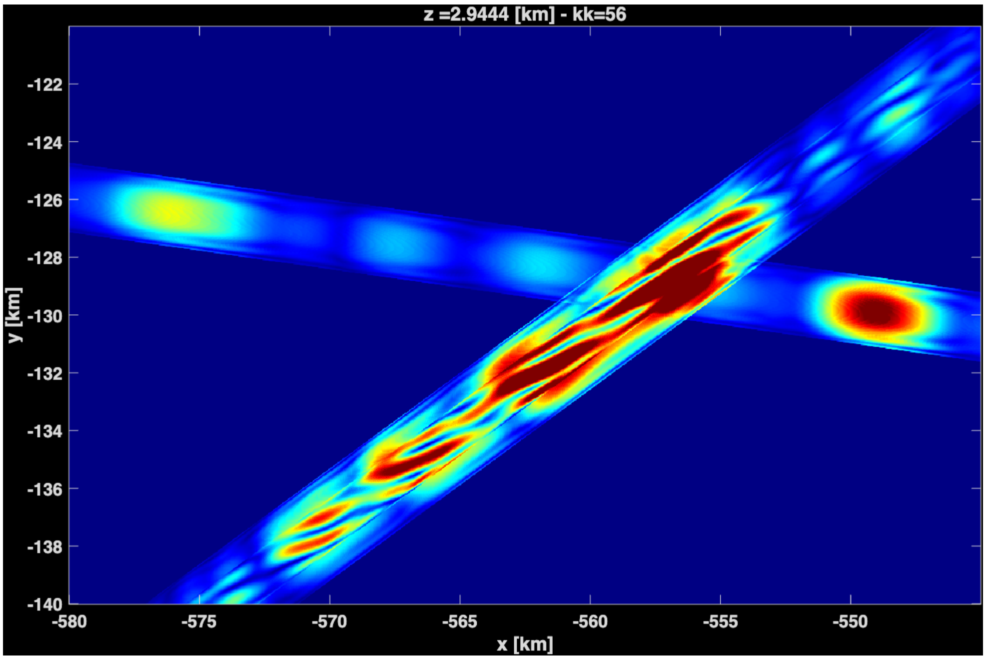

Figure 3.28: Slice of the reconstruction volume at ∼ 3 km depth. From L2C noise and phase corrected data.

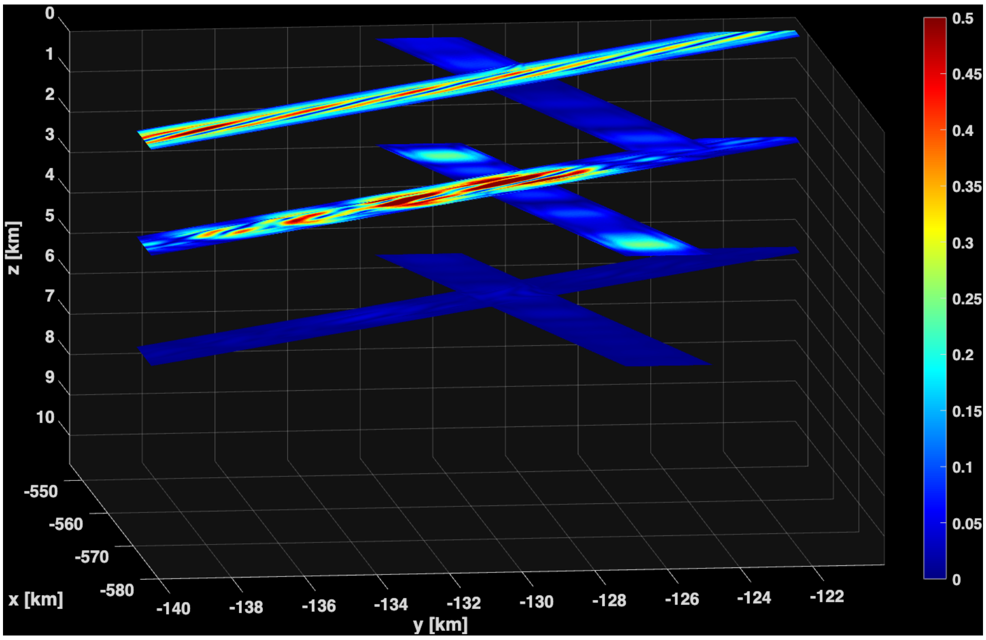

Figure 3.29: Slices of the reconstruction volume at ∼ 0, ∼ 3 and ∼ 5.5 km depth. From L2C noise and phase corrected data.

To provide a more complete picture of the three-dimensional structure of the reconstructed volume, Figure **3.30** shows a comparison between the two reconstructions. The volumetric rendering confirms all the observations made from the horizontal slices. In both reconstructions the signal energy is strongly confined to vertical columns above the orbit ground tracks, and the dominant reflector is localized at an apparent depth of approximately 3 km. Qualitatively, the L1C volume appears more diffuse, with amplitude distributed over a larger number of separated isosurface blobs; the corrected volume is visibly more compact, with fewer but better-defined structures, consistent with improved phase coherence across orbits. It is also worth noting that the isosurface blobs that appeared separated in the L1C reconstruction here resemble a unique, connected, and elongated structure.

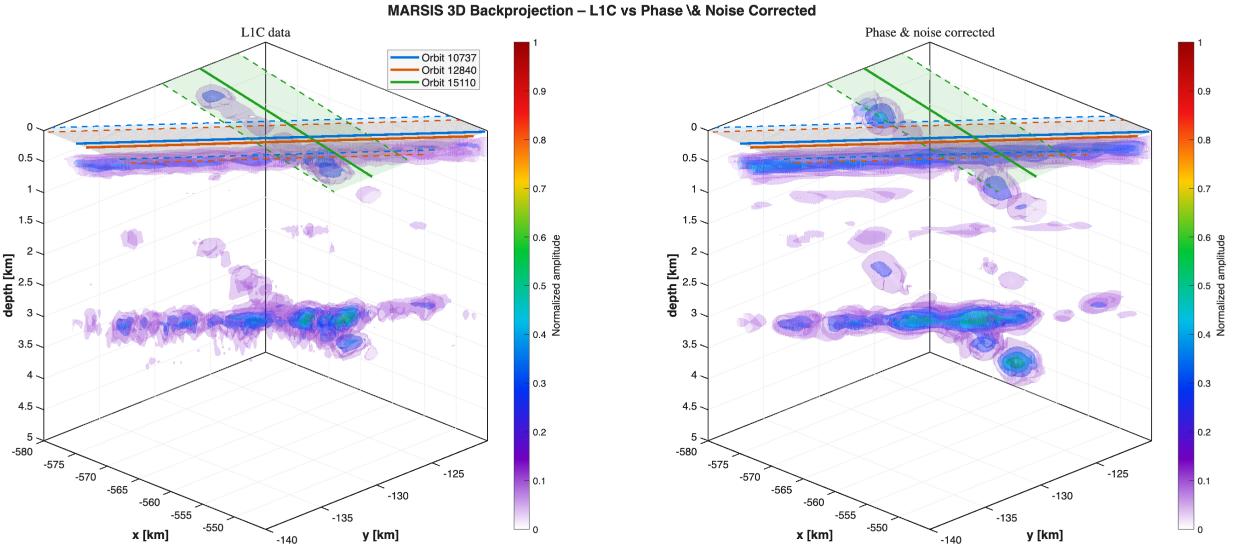

Figure 3.30: Side-by-side volume rendering of the two 3-orbits reconstructions.The visualization is based on stacked semi-transparent isosurfaces at six amplitude thresholds, ranging from 0.10 to 0.84 of the normalized maximum, with increasing opacity at higher values. The color scale encodes amplitude from violet (low) to red (high), with zero-amplitude regions left fully transparent so that the white background is visible wherever no signal is present. The Fresnel footprints of the three contributing orbits, computed from Equation **3.35**, are projected onto the surface plane ($z = 0$) as colored bands, with dashed borders indicating the $\pm r_F$ lateral extent and a solid line marking the nominal ground track.

A quantitative estimate of the vertical extent of the dominant reflector is provided by the depth profiles shown in Figure **3.31**, extracted at the pixel of maximum amplitude in each reconstruction. The left panel shows the normalized amplitude as a function of depth; the right panel shows the same profiles in decibels, relative to the respective peak, with the -3 dB and -10 dB thresholds marked for reference.

The -3 dB width of the peak is approximately 0.3–0.4 km for the L1C profile and slightly narrower for the corrected one, consistent with the tighter phase alignment after correction. This value is consistent with the range resolution expected (150 m in vacuum), broadened by the finite signal-to-noise ratio and residual phase errors. A secondary feature is visible near $z \approx 0$, corresponding to the surface echoes.

At depths greater than approximately 4 km, both profiles decay to a noise floor well below -20 dB relative to the peak, confirming that no significant coherent reflector is present below the identified structure. The depth extent of the reconstructed reflector, defined as the interval over which the amplitude exceeds the −10 dB threshold, is approximately 0.8–1.0 km. However, this value should be interpreted with caution: if the medium has a high loss tangent, as would be expected for briny water, MARSIS may lack the sensitivity to detect a reflection from the bottom of the layer, and the observed −10 dB width would then reflect only the upper boundary of the interface convolved with the system point spread function, rather than its full vertical extent.

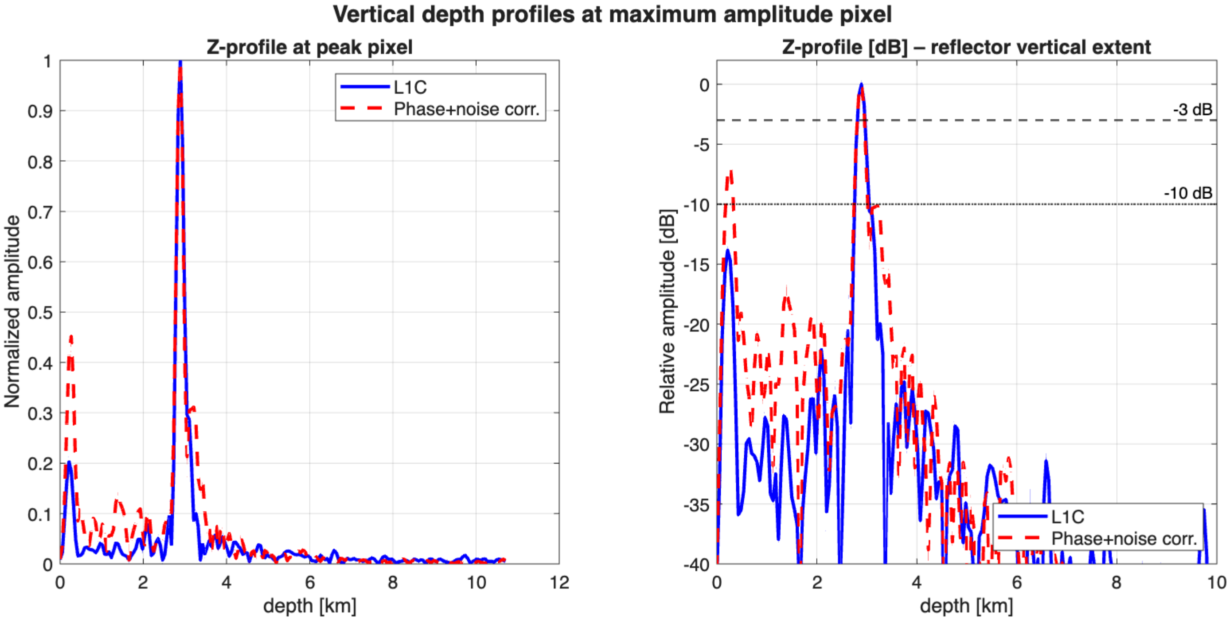

Figure 3.31: Depth profiles of the two reconstructions. The left panel shows the normalized amplitude as a function of depth; the right panel shows the same profiles in decibels, relative to the respective peak, with the −3 dB and −10 dB thresholds marked for reference.

The 27-orbit noise-corrected reconstruction was then carried out following the same procedure. To accommodate the wider spatial extent of the 27 contributing orbit tracks, the reconstruction domain was enlarged with respect to the three-orbit case, spanning:

$$x \in [-635, -520]\,\text{km}, \quad y \in [-180, -127]\,\text{km} \tag{3.39}$$

with a horizontal grid spacing of $dx = dy = 60$ m. Figure **3.32** shows the surface projection of all contributing orbit ground tracks within this extended domain. While the tracks span a wide range of azimuth angles, providing multi-directional coverage across the domain, inspection of the figure reveals that true orbit intersections are sparse: most pixels in the reconstruction domain receive contributions from at most two or three orbits simultaneously, and large portions of the domain are illuminated by a single track only. As a consequence, the angular diversity available at any given pixel is locally limited, and the reconstruction retains the characteristic imprint of the acquisition geometry: pixels near orbit crossings benefit from the most constrained spatial localization, while pixels along isolated tracks remain poorly resolved in the cross-track direction. The improvement over the three-orbit case is therefore primarily one of spatial coverage, since the reconstructed volume now extends over a larger region, rather than of angular diversity at individual pixels. Figure **3.33** shows the same stacked slices generated for the previous reconstructions, while Figure **3.34** shows the volumetric rendering of the reconstructed amplitude. In both figures the amplitude has been normalized according to two different normalization strategies. The first normalizes each orbit by its own maximum echo amplitude, as it has been done before in this analysis. This approach

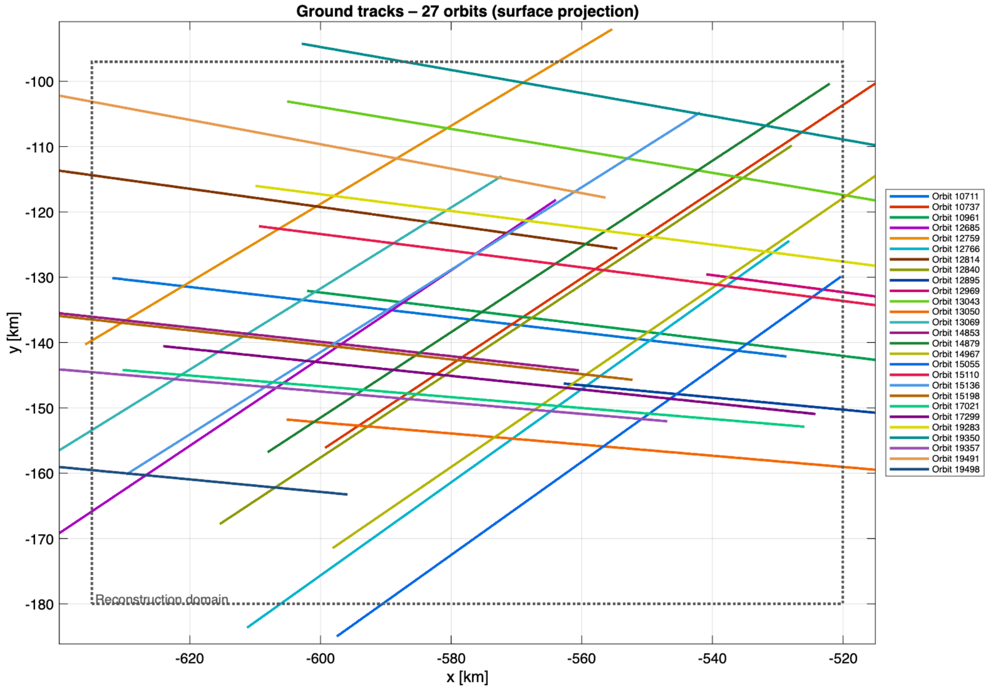

Figure 3.32: Surface projection of the ground tracks of all 27 orbits contributing to the noise-corrected 3D reconstruction, within the extended reconstruction domain $x \in [-635, -520]$ km, $y \in [-180, -97]$ km (dashed rectangle). Each orbit is assigned a distinct color for identification.

implicitly assigns equal weight to all orbits regardless of their actual signal strength, potentially allowing orbits with a higher noise floor to contribute disproportionately. The second strategy normalizes each orbit by the median of the surface echo peak amplitude. Since the surface echo is, in principle, the strongest and most stable return in every radargram, this normalization ties each orbit's amplitude scale to a physically consistent reference, effectively compensating for differences in SNR and ionospheric attenuation across orbits. Crucially, it also preserves the correct relative amplitude between the surface and any subsurface reflector: if a basal echo were to exceed the surface echo in amplitude this excess would be directly visible in the normalized volume without being suppressed by the per-orbit rescaling. In principle, the surface-normalized strategy should therefore yield a more physically interpretable reconstruction, at the cost of slightly amplifying noise in the weakest orbits.

In both cases the dominant feature remain the coherent structure localized at an apparent depth of approximately 2.9 km. The signal is no longer confined to isolated columns above individual orbit tracks, as was the case in the three-orbit reconstruction, but forms a more spatially extended and continuous structure, a direct consequence of the increased angular coverage provided by the 27-orbit aperture. Comparing the two panels, the per-orbit maximum normalization produces a somewhat more diffuse volume, with amplitude spread over a larger number of low-level blobs particularly in the $1 - 2$ km apparent depth range, while the surface-normalized volume shows marginally reduced background clutter at intermediate depths. However, the two reconstructions are qualitatively very similar.

A further element of interpretation emerges when the geometry of the reconstructed reflector is compared with the results of [1], reported in Section **3.3.3**, whose dataset comprehends the 27 orbits used in the present reconstruction. In that work, inversion of the basal echo amplitude yielded a map of relative dielectric permittivity (Figure **3.7**), with a central region of anomalously high values ($\epsilon_r \geq 15$) elongated roughly in the east-west direction, surrounded by secondary patches of elevated permittivity separated by areas of lower values. The three-dimensional structure recovered here is qualitatively consistent with this picture: the dominant reflector at $\sim 3\,\text{km}$ depth forms an elongated, connected structure whose spatial extent and orientation are broadly compatible with the high-permittivity region identified by [1].

It must be stressed, however, that this interpretation is qualitative. A quantitative comparison with the permittivity inversion of [1] would require absolute amplitude calibration and a more rigorous treatment of the three-dimensional point spread function, which is beyond the scope of the present work due to the intrinsic limitations of the problem extensively discussed in this Chapter.

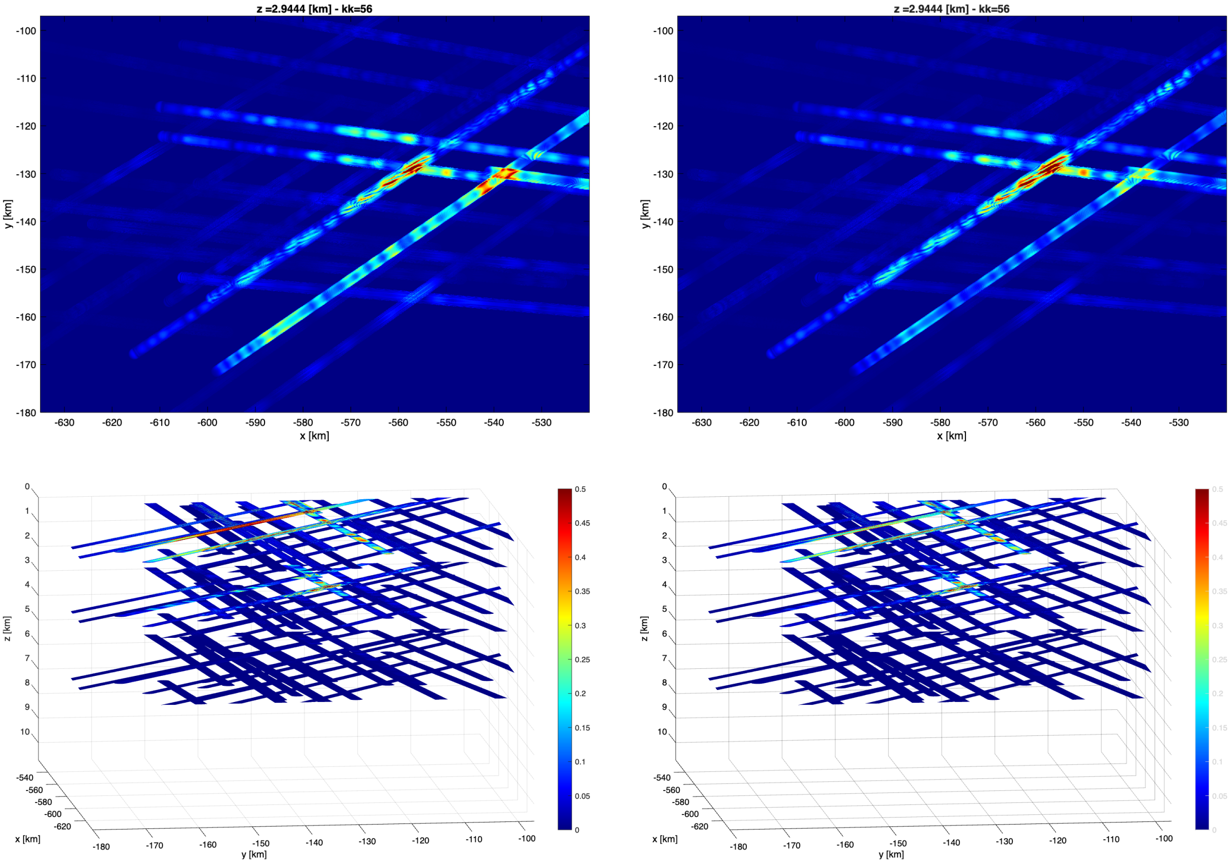

Figure 3.33: **Top**: Slice of the 27-orbits reconstruction volume at ~ 3 km depth. **Bottom**: Slices of the reconstruction volume at ~ 0, ~ 3 and ~ 5.5 km depth. **Left**: each orbit normalized by its own maximum amplitude. **Right**: each orbit normalized by the median of the surface echo peak, so that orbits with stronger ionospheric attenuation or lower SNR are upweighted relative to the dominant ones.

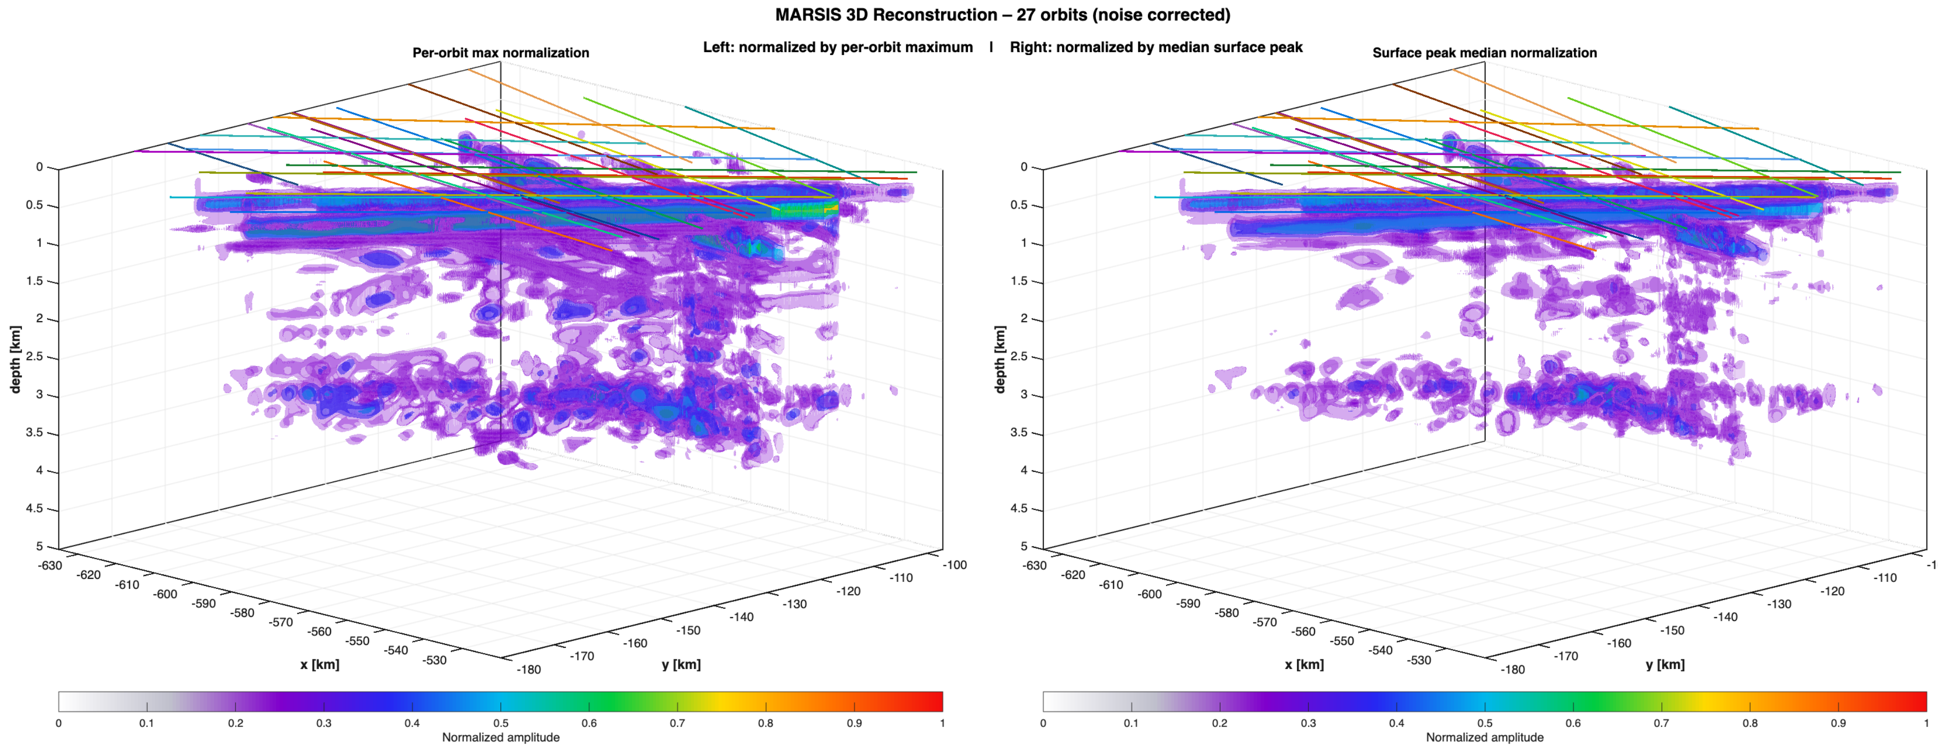

Figure 3.34: Volumetric rendering of the 27-orbit noise-corrected reconstruction under two normalization strategies. **Left**: each orbit normalized by its own maximum amplitude. **Right**: each orbit normalized by the median of the surface echo peak, so that orbits with stronger ionospheric attenuation or lower SNR are upweighted relative to the dominant ones. Both panels use identical isosurface levels (0.08 to 0.88 of the normalized maximum) and color scale. The dominant reflector at $\sim 3\,\text{km}$ depth is consistently reproduced in both cases.

Chapter 4

Ground-based radar imaging of Near-Earth Objects

This Chapter presents the application of radar techniques to the observation and characterization of Near-Earth Objects (NEOs), addressing both the instrumentation and methodology of ground-based planetary radar astronomy and the development of shape reconstruction algorithms.

The Chapter is organized as follows. We begin with an introduction to NEOs, covering their classification, orbital characteristics, scientific importance, and role in planetary defense (Section **4.1**). In Section **4.2** we describe the field of planetary radar astronomy, from its historical development to the current ground-based facilities and the continuous wave (CW) and delay-Doppler observation techniques. Section **4.3** is largely based on [155], in which we present a feasibility study for the development of a European facility for NEO radar observations. It includes a performance analysis, a description of the receiving network, and two representative observation campaigns targeting asteroids 2021 AF8 and (4660) Nereus. A detailed case study of asteroid 2005 LW3, observed with CW radar and analyzed by spectral methods for shape modeling, is presented in Section **4.4**. We then describe in Section **4.5** a progressive series of 3D shape reconstruction algorithms of increasing complexity, from a 2D convex hull model to a radial perturbation model and an adaptive vertex model with local perturbations, and present the results obtained for each approach. Finally, Section **4.6** introduces the VLBI technique and its application to asteroid radar observations, concluding with a proposed VLBI observation campaign targeting asteroid 2025 FA22 (Section **4.7**).

4.1 Introduction to Near-Earth Objects

4.1.1 NEOs classifications and orbital characteristics

Asteroids are small rocky bodies ranging in size from few meters to few hundreds of kilometers. Most of them are located in the Main Belt, which separates the terrestrial planets from the giant planets, typically between 2.1 and 3.3 AU from the Sun. Main Belt asteroids can undergo orbital modifications caused by collisions, gravitational perturbation (e.g., planetary close encounters), or non-gravitational perturbations (e.g. the Yarkovsky effect for asteroid, or outgassing for comets) that can transport them into the inner Solar System. The small bodies having orbits close to (or crossing) that of the Earth are called *near-Earth objects* (NEOs). Formally, NEOs are defined as asteroids and comets with perihelion distances $q \leq 1.3$ AU, but some of then can potentially approach Earth to within tens of millions of kilometers, and in some cases pass within the Moon's orbital radius (~ 384000 km) or even impact the planet. At the moment of writing, more than 40000 NEOs have been discovered, the majority of them with a size smaller than 100 meters but almost 1000 greater than 1 km [156].

The NEO population is of profound interest for multiple reasons spanning planetary science, astrobiology, resource utilization, and planetary defense (see Section **4.1.2**). From a practical perspective, NEOs represent both hazard and opportunity. Impact events, while rare on human timescales, could cause enormous damage to the planet: a 140 meter diameter asteroid impact releases energy equivalent to hundreds of megatons of TNT, capable of regional devastation, while kilometer-scale impacts could trigger global climate disruption. In contrast, NEOs are the most accessible extraterrestrial resources, which require slower velocity changes (Δv) for reaching the lunar surface in many cases and potentially contain water, metals and other valuable materials for future space exploration and industrial development.

A sketch of near-Earth asteroids (NEAs) and their orbits is shown in Figure **4.1**. NEOs are subdivided into four principal orbital classes based on their semi-major axis a, perihelion distance q, and aphelion distance Q, reflecting different dynamical origins and Earth encounter geometries [157, 158]:

- Atira asteroids (or Interior-Earth Objects, IEOs) orbit inside Earth's orbit. Specifically, they have an aphelion distance $Q < 0.983$ AU and a semi-major axis $a < 1.0$ AU. Atira asteroids are by far the least numerous group of NEOs, since their observation is challenging due to the small solar elongation and their daytime only visibility.

- Aten asteroids have a semi-major axis $a < 1.0$ AU, and an aphelion distance $Q > 0.983$ AU. This class has been named after 2062 Aten, discovered on 7 January 1976 by American astronomer Eleanor Helin at Palomar Observatory. As of January 2025, 2,966 Atens have been discovered.

- The Apollo asteroids are a group of near-Earth asteroids named after 1862 Apollo, discovered by German astronomer Karl Reinmuth in the 1930s. Their semi-major axis is $a \geq 1.0$ AU and perihelion distance q ≤ 1.017 AU, less then Earth's aphelion. Currently, they are the most numerous NEO class, with more than 21000 objects classified. Among them, more than 2000 have been identified as Potentially Hazardous Asteroids (PHA), which are asteroids that make close approaches to Earth and with a size large enough to cause significant regional damage in the event of impact [159].

- Amor asteroids, named after the object 1221 Amor, have a perihelion distance $1.017 < q \leq 1.3$ AU. Hence, these objects do not cross Earth's orbit. Amor asteroids occupy a dynamically unstable region of orbital phase space characterized by strong gravitational perturbations from both Earth and Mars. Numerical integration studies of Amor orbital evolution reveal that the majority of these objects will transition to Earth-crossing (Apollo or Aten) orbits within timescales of 10^5 to 10^6 years through several mechanisms [160, 161]. Currently, more than 15000 Amor asteroids have been detected.

While the four classes described above refer exclusively to near-Earth asteroids, the NEO population also includes near-Earth comets (NECs), defined as comets with a perihelion distance $q \leq 1.3$ AU and an orbital period $P \leq 200$ years. NECs are volatile-rich bodies subject to non-gravitational forces such as outgassing, dynamically and physically distinct from NEAs. Currently, ~ 4000 comets have been discovered [162], representing a small fraction of the total NEO population. As mentioned before, the orbital evolutions of NEOs is dominated by gravitational perturbations from the terrestrial planets (primarily Earth, Venus, and Mars) and, to a lesser extent, Jupiter [163]. Close approaches to Earth can be characterized by several geometric parameters:

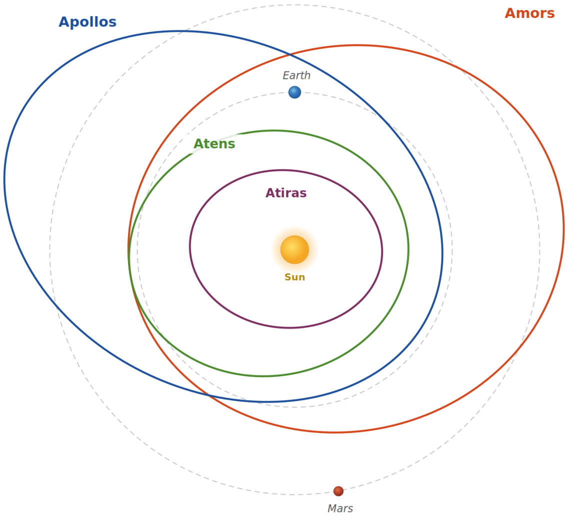

Figure 4.1: Sample schematic orbits of objects in the four traditional asteroidal NEO classes

- the *Minimum Orbit Intersection Distance* (MOID) is defined as the minimum possible distance between the orbits of the NEO and Earth, neglecting planetary masses. This parameter is a lower bound on close approach distance. Objects with MOID < 0.05 AU are classified as PHA if they also exceed 140 m in diameter.

- The *close approach distance* is defined as the actual geocentric distance during a specific encounter, which depends on the relative positions of the NEO and Earth along their orbits at the time of conjunction. Close approaches can range from millions of kilometers to, in rare cases, within Earth's gravitational sphere of influence (~ 930000 km).

- The velocity of the NEO relative to Earth at close approach is determined by the vector difference between the NEO's heliocentric velocity and Earth's orbital velocity. Typical values range from ~ 5 km/s (for low-inclination, low-eccentricity objects with orbits nearly matching Earth's) to > 30 km/s (for high-inclination or highly eccentric objects).

The geometry of Earth encounters is crucial for adopting a ground-based observational approach. Radar observations require close approaches, typically within $\sim 0.05 - 0.15$ AU, to achieve sufficient SNRs with available planetary radar systems. Optical observations can detect objects at greater distances, albeit with reduced spatial resolution and photometric precision.

Statistically, the number of NEOs can be expressed as a diameter distribution [164]:

$$N(> D) \propto D^{-\alpha} \tag{4.1}$$

where N(>D) is the cumulative number of objects larger than diameter D, and the slope α varies with size range. For D > 1 km, $\alpha \approx 2.3 - 2.5$; for smaller objects, the distribution may steepen

[165].

Compositionally, NEOs exhibit diversity reflecting their origins from different source regions in the Main Asteroid Belt. Spectroscopic surveys in visible and near-infrared wavelengths, combined with albedo measurements, reveal three principal compositional families that account for $\sim 90\%$ of the NEO population [166, 167]:

- C-type (Carbonaceous) asteroids are characterized by very dark surfaces with geometric albedos typically in the range 0.03-0.10, making them among the least reflective objects in the Solar System. The low albedo is due to the presence of opaque minerals such as magnetite, carbon-rich organics, and fine-grained phyllosilicates that efficiently absorb incident sunlight across visible wavelengths. The presence of hydrated minerals suggest the formation or the alteration of these objects in water-rich environments. C-type asteroids are interpreted as primitive bodies preserving relatively unmodified material from the early Solar System. C-types are dominant in the outer Main Belt (beyond ~ 2.7 AU), where lower formation temperatures allowed volatile retention. However, they constitute $\sim 20-30\%$ of the NEO population, having been transported to Earth-crossing orbits through resonant ejection mechanisms. C-type asteroids generally exhibit moderate radar reflectivity but high circular polarization ratios ($\mu_C = \mathrm{SC/OC} > 0.5$), where SC (same-sense circular) and OC (opposite-sense circular) refer to the polarization state of the received echo relative to the transmitted signal (see Section **1.2.2**). A high μ_C indicates significant wavelength-scale surface roughness and is the primary radar observable for the physical characterization of NEA surfaces and sub-surfaces [168]. The presence of hydrated minerals may also affect radar penetration depth and volume scattering properties.

- S-type (Silicaceous) asteroids exhibit moderate geometric albedos in the range $0.10-0.25$, due to the presence of silicate minerals, such as olivine ($(\mathrm{MgFe})_2\mathrm{SiO}_4$) and pyroxene ($(\mathrm{MgFe})\mathrm{SiO}_3$), that are also the dominant materials in the Earth's mantle. In addition to that, many S-types also show evidence for metallic iron-nickel in their surface regolith. S-types are interpreted as partially differentiated bodies that experienced a form of thermal heating sufficient to dehydrate volatile-bearing minerals and, in some cases, cause partial melting and metal-silicate segregation. They are the dominant population in the inner Main Belt (up to ~ 2.5 AU), where higher formation temperatures could lead to volatile loss. They constitute $\sim 40-50\%$ of the known NEO population, making them the most numerous compositional class among Earth-approaching objects. For radar observations, S-type asteroids generally exhibit moderate to high radar reflectivity (depending on their metal content) and low to moderate circular polarization ratios ($\mathrm{SC/OC} = 0.1-0.4$), indicating surfaces with a mixture of smooth facets and wavelength-scale roughness.

- M-type (Metallic) asteroids exhibit high geometric albedos ($0.10-0.60$), due to the efficient reflection of sunlight by metallic surfaces. Even if they are mostly composed of metals, some of them show evidence for minor silicate components. M-type asteroids are interpreted as exposed metallic cores or core fragments from differentiated parent bodies that underwent complete melting and metal-silicate segregation early in Solar System history, followed by catastrophic collisional disruption that stripped away the silicate mantle and crust. M-types are relatively rare among NEOs, constituting $\sim 5-10\%$ of the population. They are most common in the middle Main Belt ($\sim 2.4-2.7$ AU), with the largest concentration in the vicinity of the well-known 16 Psyche. For radar observations, they exhibit extremely high radar reflectivity and very low circular polarization ratios (SC/OC < 0.1), indicating smooth surfaces at radar wavelengths dominated by coherent specular reflection from metallic facets. The metallic content also results in low radar penetration depth ($<$ 1 meter). These radar properties make M-types asteroids uniquely identifiable, since high optical albedos are not sufficient for their characterization: some S-type asteroids can also exhibit similar values.

Most NEOs have orbital periods ranging from ~ 2 hours to several days, determined from optical lightcurve observations. A significant fraction ($\sim 15-20\%$) of NEOs are suspected or confirmed binary or multiple systems, where a primary body is accompanied by one or more smaller satellites [169]. Binary formation mechanisms likely include rotational fission, tidal disruption during close planetary encounters, and collisional fragmentation followed by gravitational re-accumulation. Radar observations are particularly effective at detecting and characterizing binary systems through analysis of echo power distributions and Doppler spectra, which can reveal separated components unresolved in optical observations. The case study of 2005 LW3 presented in Section **4.4** is a clear example of these detections.

4.1.2 Scientific importance and planetary defense

Asteroids represent some of the most primitive material accessible to scientific investigation: aggregations of dust, ice, and rocky fragments that condensed from the solar nebula approximately 4.5 billion years ago but never experienced the high temperatures and pressures of planetary differentiation. Many asteroids preserve relatively pristine samples of nebular material, offering direct constraints on the composition, temperature, and dynamics of the protoplanetary disk. Also, the size, orbital, and compositional distributions of NEOs preserve information about the collisional and dynamical evolution of the asteroid belt, as well as transport mechanisms and impact rates on terrestrial planets.

Carbonaceous asteroids, given the presence of volatile compounds, likely formed beyond the "snow line", that is the heliocentric distance (approximately 2.7 AU in the current epoch, but variable during disk evolution) where water ice could condense from the nebula. The detection of water-bearing material indicates that liquid water existed in the asteroid belt from the primordial phase of the planets formation. These compounds may have synthesized complex organic molecules relevant to the origin of life.

Silicate asteroids formed interior to the snow line, since they do not possess volatile components. Spectral matches between S-type asteroids and ordinary chondrite meteorites establish a genetic link between these populations, confirming that much of the meteorite flux to Earth originates from Main Belt asteroids perturbed into planet-crossing orbits [170].

Metallic asteroids likely represent fragments of differentiated planetesimals that experienced core-matle separation [171]. The existence of differentiated asteroids demonstrates that planetesimal accretion and melting occurred within the first few million years of Solar System history, providing chronological constraints on planetary formation timescales.

The recognition that asteroid impacts pose a quantifiable threat to human civilization emerged from geological studies of the Cretaceous-Paleogene impact 66 million years ago, which originated the ~ 180 km Chicxulub crater in Mexico and triggered mass extinctions [172]. The impact hazard is quantified through statistical analysis of NEO population, collision cross-sections, and impact energies. Current estimates indicate [165, 173]:

- D ∼10 km impacts occur every $\sim 10^8$ years, causing global catastrophic effects (mass extinction, climate disruption);
- D ∼ 1 km impacts occur every $\sim 10^6$ years, causing regional-to-continental devastation and global climate perturbation;
- D ∼ 140 m impacts occur every $\sim 10^4$ years, causing regional devastation (area $\sim 10^4$ km^2) with an energy release estimated amount of ∼10-100 megaton;
- D ∼ 30 m impacts (Chelyabinsk-class) occur every ∼100 years, causing local damage with an airburst energy ∼0.5 megatons;
- D ∼ 5-10 m impacts occur annually. The resulting catastrophic fragmentation and rapid deceleration deposits kinetic energy into the atmosphere as a high-altitude airburst, producing luminous bolide phenomena visible over areas of thousands of square kilometers [174].

The establishment of systematic NEO surveys with the primary aim of detecting $\geq 90\%$ of NEOs with $D \geq 1$ km (in 1998, subsequently extended to the 140-m threshold) has dramatically increased discovery rates and improved orbital knowledge for potentially hazardous objects.

4.2 Planetary Radar Astronomy

4.2.1 Historical development

Planetary radar astronomy is a field of science at the intersection of planetology, radio astronomy, and radar engineering [175]. The development of this field began after World War II, when radar engineers recognized that the powerful transmitters and sensitive receivers developed for military applications might be capable of detecting echoes from celestial bodies. On January 10, 1946, the Moon became the first astronomical target ever successfully detected, using a radar operating at 111.5 MHz. The first successful radar detection of a planet, Venus, was achieved nearly simultaneously by several research teams in 1961, representing a milestone that validated planetary radar as a tool for accurate distance measurement within the Solar System.

The extension of planetary radar to small bodies - asteroids and comets - became feasible as transmitter powers and receiver sensitivities improved through the 1960s. The first radar detection of an asteroid was achieved at Goldstone in June 1968, when echoes were detected from 1566 Icarus during a remarkably close approach to within 0.04 AU (∼ 6 million km) of Earth [176]. Icarus, an Apollo-type near-Earth asteroid with a highly eccentric orbit (perihelion 0.19 AU, aphelion 1.97 AU) and a diameter of approximately 1.3 km, presented a favorable target due to its proximity and relatively large size. The successful detection, achieved with a transmitter power of ∼ 450kW at X-band (8.51 GHz) using the 64-meter DSS-14 antenna at Goldstone (later enlarged to 70 m between 1987 and 1988, as referred to throughout this work), demonstrated that near-Earth asteroids could be characterized by radar during close approaches, establishing the feasibility of systematic NEO radar observations.

Early asteroid radar detections focused primarily on measuring echo power and Doppler bandwidth to estimate sizes and rotation periods. These measurements, while limited compared to modern imaging capabilities, provided valuable constraints on asteroid physical properties unavailable from optical observations alone. The radar-derived diameter estimates were free from the albedo-size degeneracy that is present in optical photometry, while Doppler bandwidth measurements provided first-order constraints on rotation rates and pole orientations.

The first delay-Doppler images of an asteroid, enabling detailed shape reconstruction rather than merely detection and size estimation, were obtained of 4179 Toutatis during its December 1992 close approach to within 0.024 AU (~ 3.6 million km) of Earth [177]. Observations at both Arecibo and Goldstone over multiple days produced delay-Doppler images with spatial resolutions of approximately 125 m $\times$ 125 m, revealing Toutatis as an elongated, bifurcated contact binary asteroid ~ 4.0 km $\times$ 2.4 km $\times$ 1.9 km, exhibiting a complex non-principal axis rotation state with two independent rotation periods of ~ 5.4 and ~ 7.3 days. This tumbling rotation, caused by the object's highly asymmetric shape and non-principal-axis angular momentum vector, produces a complex pattern of echo power and Doppler variations that enabled detailed physical modeling.

The physical modeling methodology developed for Toutatis-iterative shape refinement through comparison of observed and synthetic delay-Doppler images was subsequently applied to dozens of NEOs, becoming the standard approach for asteroid shape reconstruction from delay-Doppler radar data [44, 178].

4.2.2 Ground-based radar facilities

Ground-based planetary radar observations require high-power transmitters and large antennas. As shown in Eq. **1.68**, the received echo power scales as $P_r \propto R^{-4}$ due to the geometric spreading of the signal over the two-way propagation path. This dependence on distance makes the detection of small or distant targets extremely challenging, and only a handful of facilities worldwide have ever been capable of conducting planetary radar observations of near-Earth objects.

The Arecibo Observatory in Puerto Rico [179] operated the world's most powerful planetary radar system from 1963 until its catastrophic collapse on 1 December 2020. The facility combined a 305-meter reflector with a megawatt-class transmitter, and the dual-reflector system installed in 1997 significantly improved beam quality and aperture efficiency. It enabled detection and imaging of NEOs as small as $\sim$50 m at distances exceeding 0.05 AU. The loss of this facility eliminated approximately 60–70% of global planetary radar capacity.

The DSS-14, located at the Goldstone Deep Space Communications Complex (GDSCC) [180] in California's Mojave Desert hosts the Goldstone Solar System Radar (GSSR), which is the world's only dedicated planetary radar still in operation. Unlike Arecibo's fixed geometry, Goldstone employs a fully steerable 70-meter antenna providing all-sky access. Goldstone's X-band frequency provides approximately three times better range resolution than Arecibo's S-band for equivalent bandwidth, though at the cost of increased atmospheric attenuation and reduced aperture efficiency. The facility operates in both monostatic (same transmitting/receiving antenna) and bistatic (separate transmitting/receiving antennas) configurations.

DSS-63, the 70-meter antenna at NASA's Madrid Deep Space Communications Complex near Robledo de Chavela, Spain, is one of three 70-meter Deep Space Network (DSN) stations providing continuous spacecraft coverage. Primarily dedicated to spacecraft communication, DSS-63 operates in the C-band, specifically in its upper frequency portion, conventionally referred to as C-high ($6.0 - 8.0$ GHz), transmitting capability with power up to 20 kW. This is significantly lower than DSS-14's 450 kW planetary radar transmitter, reflecting DSS-63's dedication to spacecraft uplink rather than planetary target illumination. The antenna has occasionally been used for bistatic planetary radar observations, transmitting while other facilities receive, though such operations are rare due to competing spacecraft communication priorities.

DSS-43, the 70-meter antenna at NASA's Canberra Deep Space Communication Complex (CDSCC) near Tidbinbilla, Australia, completes the DSN's three-station global network alongside Goldstone and Madrid [181]. Like DSS-63, DSS-43 is primarily dedicated to spacecraft communication, but is equipped with an C-High transmitter capable of up to 80 kW. While this exceeds DSS-63's 20 kW capability, it remains well below the 450 kW of DSS-14. DSS-43's southern hemisphere location provides complementary sky coverage to Goldstone, making it a valuable asset for bistatic campaigns targeting objects with low or negative declinations.

Historically, beyond Arecibo and the currently operational Goldstone system, one of the few planetary radar transmitters was the 70-meter RT-70 antenna at the Evpatoria Deep Space Center in Crimea. During the 1960s-1980s Soviet space program, this facility conducted pioneering planetary radar observations of Venus and Mars. The RT-70 was equipped with a 5010 MHz transmitter capable of delivering up to 200 kW of continuous power, which was also utilized for deep-space spacecraft communication. While the facility demonstrated significant planetary radar capabilities - including NEO observations ([182]) - and was subsequently used for radio astronomy and spacecraft tracking [183], it is no longer operational. The Evpatoria complex seems to have suffered severe damage from military strikes during the Russo-Ukrainian War, effectively ending its operational life ([184]). A summary of the main characteristics of these facilities is given in Table **4.1**. Beyond the primary 70-meter facilities, a number of smaller 34-meter DSN antennas have been employed as radar transmitters in selected experiments, albeit with performance limited by their comparatively small aperture and correspondingly lower antenna gain. At the Canberra Deep Space Communication Complex (CDSCC), the 34-meter antennas DSS-34, DSS-35, and DSS-36 have been used in bistatic configurations, transmitting toward NEO targets with reception handled by smaller facilities such as the 12-meter antennas of the University of Tasmania [185]. Similarly, DSS-13 at Goldstone has been utilized as a transmitter in planetary radar experiments, particularly when DSS-14 is unavailable [186].

Several initiatives aim to expand or restore planetary radar capabilities:

- Next Generation Arecibo Telescope (NGAT)is a proposed 300-meter phased array design to replace the collapsed Arecibo facility, potentially offering improved sensitivity and sky coverage [187];
- currently several upgrades to 70-meter DSN antennas at Canberra (Australia), and Madrid (Spain) with higher transmitter powers are under discussion;
- the idea of combining multiple smaller transmitters to achieve high EIRP through coherent phased-array techniques (Distributed Transmitter Arrays) is under study. Rather than relying on a single megawatt-class transmitter, distributed architectures would phase-synchronize multiple moderate-power transmitters (e.g., 10-50 kW each) across apertures of hundreds of meters, coherently combining their outputs to achieve effective radiated powers comparable to traditional facilities [188].

While high-power transmitters are rare, numerous large radio telescopes worldwide can serve as sensitive receiving stations for bistatic planetary radar observations. These facilities, when paired with distant transmitters, enable radar characterization of NEOs through internationally coordinated campaigns. Major receiving facilities recently employed for planetary radar include the following:

Table 4.1: Comparison between major radar transmitter systems: Arecibo (final configuration, 1997–2020), Goldstone DSS-14, Canberra DSS-43, Evpatoria RT-70, and Madrid DSS-63.

	Arecibo	**Goldstone DSS-14**	**Canberra DSS-43**	**Evpatoria RT-70**	**Madrid DSS-63**
Diameter (m)	305	70	70	70	70
Frequency (MHz)	2380 (S-band)	8560 (X-band)	7167 (C-high)	5010 (C-band)	7167 (C-high)
Wavelength (cm)	12.6	3.5	4.2	6.0	4.2
Tx power (kW)	1000	450	80	200	20
Gain (dBi)	∼73	∼74	∼74	∼69	∼74
EIRP (TW)	∼20	∼12.6	∼2.0	∼1.7	∼0.5
Operational	No	Yes	Yes	No	Yes

*DSS-43 and DSS-63 are used exclusively as transmitters in bistatic radar configurations for NEO observations.

- Green Bank Telescope (USA): 100 m offset Gregorian design;
- Effelsberg Radio Telescope (Germany): 100 m;
- Jodrell Bank Lovell Telescope (UK): 76 m;
- Sardinia Radio Telescope – SRT (Italy): 64 m;
- Parkes Radio Telescope (Australia): 64 m;
- Medicina “G. Grueff” Radio Telescope (Italy): 32 m;
- Irbene Radio Telescope (Latvia): 32 m;
- Ceduna Radio Telescope (Australia): 30 m;
- Very Large Array – VLA (USA): 27 × 25 m;
- Very Long Baseline Array – VLBA (USA): 10 × 25 m;
- Australia Telescope Compact Array – ATCA (Australia): 6 × 22 m.

The loss of Arecibo has caused a fundamental shift in planetary radar observational strategies. The pre-2020 paradigm relied on one or two dominant facilities (Arecibo and Goldstone) conducting primarily monostatic observations with complete control over transmitter timing, waveform, and receiver synchronization. The post-Arecibo era increasingly exploited distributed bistatic observations: relatively lower-power and/or smaller size transmitters paired with large receiving apertures (e.g. Effelsberg, Green Bank) located at continental distances. The bistatic approach has both challenges and opportunities. Challenges include complex time synchronization across facilities, the difficulty of coordinating schedules between independent telescopes and the need to find a simultaneous window of mutual target visibility for both the transmitter and the receivers. Opportunities include enhanced geometric coverage through multiple simultaneous receiving stations, elimination of transmitter-receiver isolation problems, potential for bistatic scattering geometries that reveal surface properties inaccessible to monostatic observations and the ability to perform uninterrupted, coherent signal integration over much longer periods.

The development of signal processing and shape reconstruction methodologies capable of extracting maximum information from such distributed, lower-sensitivity observations represents an essential adaptation to the current constrained planetary radar scenario.

4.2.3 CW and delay-Doppler techniques

Ground-based planetary radar observations employ two fundamentally different transmission and signal processing strategies: continuous-wave transmission (Section **2.2.1**) or Delay-Doppler setting, characterized by a transmitted modulated waveforms that allow simultaneous measurement of range (time delay) and radial velocity (Doppler frequency).

CW radar transmits an unmodulated constant-amplitude sinusoidal carrier, as in Equation **2.34**. While the simplest implementation operates at a fixed frequency, in practice the transmitted frequency can be continuously adjusted to pre-compensate for the predicted Doppler shift at the receiver, a strategy routinely adopted at Goldstone to keep the echo centered within the receiver band. Echoes from the target are received simultaneously from all surface elements, with each element contributing a Doppler-shifted frequency component determined by its line-of-sight velocity relative to the radar. This produces the echo spectral broadening. However, the entire echo spectrum is also subject to a much larger bulk Doppler shift caused by the overall motion of the asteroid's center of mass relative to the observer (topocentric velocity). For typical NEO encounter geometries with relative velocities on the order of 10 km/s and radar frequencies of several GHz, this overall Doppler shift can reach up to several hundred kHz.

The received signal is a superposition of echoes from all illuminated surface elements, each contributing a frequency component proportional to its radial velocity. For a rotating asteroid illuminated broadside, surface elements on the limbs approaching and receding from the radar contribute the maximum positive and negative Doppler shifts relative to the asteroid's center of mass, respectively. Moreover, the sub-radar point (zero line-of-sight velocity) contributes a component exactly at the bulk Doppler-shifted frequency of the center of mass.

The resulting echo power spectrum, obtained through Fourier transformation of the time-domain signal, exhibits a characteristic shape reflecting the target's rotation state and projected shape. More complex target shapes produce different spectral signatures in the echo profile. Elongated objects exhibit broader bandwidths if rotating about their short axis, while irregular shapes or non-principal-axis rotation create time-varying spectral features. Gravitationally bound binary systems produce spectra with multiple distinct echo bandwidth features corresponding to the rotation of the primary and secondary bodies.

CW transmission is still widely used due to several advantages that it carries. First of all, CW transmission requires only constant-amplitude RF generation without complex pulse modulation or timing circuits, reducing transmitter complexity and cost. Continuous transmission (corresponding to a 100% duty cycle) maximizes the average transmitted power for a given peak-power capability, thereby enhancing sensitivity in power-limited systems. The absence of range gating allows integration of echoes from the entire target simultaneously, potentially improving SNR for distant or small objects compared to range-resolved techniques that divide echo power among many range cells. Finally, echo power variations at the observer's location (due to changing aspect angle and/or radar albedo) directly yield rotation periods through Fourier analysis techniques.

On the other hand, CW observations provide no information about the range: all echoes arrive simultaneously, preventing separation of front and back surface contributions. The Doppler spectrum maps only the line-of-sight velocity, making it impossible to distinguish between rotation poles oriented toward +Z or -Z (where Z is perpendicular to the plane-of-sky). This North-South ambiguity can be resolved by performing observations at multiple epochs with different viewing geometries. Also, while binary systems produce characteristic spectral features (a narrow spike from the secondary superimposed on a broad primary echo), determining mutual orbit geometry and satellite properties is more challenging than with delay-Doppler data.

Delay-Doppler radar transmits long continuous waveforms utilizing the pulse compression techniques, typically employing Binary Phase Coding (BPC) or linear frequency modulation (LFM/chirp) to achieve fine range resolution while maintaining a high average transmitted power without exceeding the transmitting power and bandwidth limitations. Range resolution is given by Equation **2.42**. Simultaneously, the Doppler frequency shift of each range-resolved echo cell is measured by Fourier analysis of the time series of echoes at that range, providing Doppler resolution:

$$\Delta f_{\text{Doppler}} = \frac{1}{T_{\text{coh}}} \tag{4.2}$$

where T_{coh} is the coherent integration time. This yields a two-dimensional delay-Doppler image in which each pixel represents echo power from a specific combination of range (distance from radar) and Doppler frequency (line-of-sight velocity).

The standard approach to asteroid shape reconstruction from radar data employs delay-Doppler images acquired over multiple rotation cycles and multiple observation epochs with different viewing geometries when possible [44, 178, 189]. This technique can be summarized in the following steps:

- define an initial simple shape model (an ellipsoid or simple convex polyhedron) approximating the target's overall dimensions, inferred from optical lightcurve and radar bandwidth;
- for a given shape, rotation state (pole direction, period), and viewing geometry, compute synthetic delay-Doppler images by raytracing: for each surface facet, calculate delay and Doppler, sum contributions from all facets mapping to each pixel;
- compute residuals between observed and synthetic images;
- adjust vertex positions of the shape model to minimize residuals;
- simultaneously optimize rotation period, pole direction, and initial rotation phase to match the time evolution of images.

A comparison between CW radar and Delay-Doppler radar techniques is reported in Table **4.2**.

Table 4.2: Comparison between Continuous Wave (CW) radar and Delay-Doppler radar techniques in planetary radar astronomy.

Characteristic	**CW Radar**	**Delay–Doppler Radar**
Transmitted waveform	Continuous unmodulated (CW)	Continuous, phase-coded (PSK) or frequency-ramped (chirp)
Range resolution	None	7.5–300 m (bandwidth dependent)
Doppler resolution	High (∼0.01–0.1 Hz)	Moderate (∼0.1–1 Hz)
Hardware complexity	Low	High (timing, modulation, compression)
Echo information	1D spectrum (Doppler only)	2D image (range × Doppler)
Shape reconstruction	Difficult, requires rotation period coverage	Standard, single epoch sufficient
Rotation period	Direct from spectrum	From image time series
Binary detection	Possible from spectral features	Clear spatial separation
Sensitivity (SNR)	High (full target integration)	Low (distributed over pixels)

4.3 Toward a European Facility for NEO radar observations

As mentioned in the previous sections, the collapse of the Arecibo Observatory in 2020 led to a significant void in planetary radar observations. This situation required a reassessment of existing and potential assets worldwide, specifically in Europe, since we lack a dedicated high-power transmitter and we rely only on NASA/JPL transmitting assets. Recognizing both the scientific and strategic dimensions of this gap, the European Space Agency issued in 2019, even before the Arecibo incident, a call for tenders for the project "NEO Observation Concepts for Radar Systems" (ESA contract SSA P3-NEO-XXII), which successfully concluded in 2022 [190]. The project aimed to define the functional requirements that a radar system must meet for NEO observations from both a scientific and a planetary defense perspective, to evaluate the European assets that might be employed or upgraded for such activities, and to carry out test radar observation campaigns demonstrating the practical potential of existing European infrastructure. The consortium awarded the ESA tender was led by SpaceDyS and included the Italian National Institute for Astrophysics (INAF), and the University of Helsinki, with the essential contribution of NASA's Jet Propulsion Laboratory (JPL), the Goldstone Deep Space Communication Complex (GDSCC) and the Madrid Deep Space Communications Complex (MDSCC) for the execution of the observation campaigns. In the wake of this project, these radar campaigns and research activities still continue.

The work described in this Section was carried out within this framework and is extensively presented in [155], of which the author is a co-author. The author's contributions included the development of signal processing software and spectral analysis pipelines, the implementation of shape reconstruction algorithms, and participation in the observation planning and data reduction for the 2006 WB and 2025 FA22 campaigns. The observations and results concerning 2021 AF8 and (4660) Nereus are presented as case studies to which the author contributed within this broader collaborative framework. The primary original contribution of this thesis is the complete analysis of asteroid 2005 LW3, from data reduction and spectral analysis to the full shape reconstruction pipeline, which was carried out entirely by the author and is described in detail in Section **4.4** and **4.5**. In particular, given the innovative use of Very Long Baseline Interferometry (VLBI) in the study of 2025 FA22, a dedicated Section is provided for that campaign.

4.3.1 Performance Analysis and System Requirements

The first task led by INAF within the ESA project consisted of a systematic quantitative analysis aimed at deriving the functional requirements for a European planetary radar system and evaluating the expected scientific return as a function of system parameters. This analysis covered transmitted and received frequency bands, pointing and tracking accuracy, signal waveform, polarization, sensitivity, measurement resolution, readiness and accuracy requirements. The results informed both the design of the test observation campaigns and the longer-term planning for a permanent European planetary radar facility.

Custom simulation software was developed in MATLAB to evaluate the expected detection rate (the number of NEOs detectable per year with SNR exceeding a given threshold) as a function of transmitted power, for four frequency bands: L (1.333 GHz), S (2.38 GHz), X (8.56 GHz), and Ka (34.0 GHz). These bands correspond to historically employed planetary and space surveillance radar systems (Tracking and Imaging Radar, TIRA, a space surveillance radar primarily designed for space debris and artificial satellite tracking) and to future high-frequency options. Hence, the technical parameters used were taken by those of real antennas, when existing, otherwise we considered plausible values from literature. Table **4.3** shows the main parameters of the sensors used in the simulations: radio frequency band with the operating frequency, antenna size, noise temperature $\mathrm{T_{AMW}}$ due to the antenna and its microwave components (without atmosphere noise temperature), and antenna gain. Atmospheric effects, such as attenuation and noise temperature contribution, were modeled using a standard atmospheric model applied [191] to four weather scenarios derived from long-term statistics measured at the Madrid DSN site, characterized by their cumulative distribution value CD [192]. For example, CD $=$ 0.25 corresponds to average clear weather conditions (25% of days have better atmospheric transparency), while CD $=$ 0.90 corresponds to very cloudy but rain-free conditions, indicating that for 90% of the year, the site experiences weather conditions equal to or more favorable than this threshold.

The SNR of the radar echo at the time of close approach was estimated using the bistatic radar equation. Starting from Eq. **1.69** and assuming that the mean noise background is stable, we define the noise standard deviation as:

$$\Delta P_N = \frac{kT_{\mathrm{sys}}\sqrt{B}}{\sqrt{\Delta t}} \tag{4.3}$$

where Δt is the integration time, k is the Boltzmann constant, T_{sys} is the system temperature, and B is the bandwidth in Hz. For a spherical object, the radar cross section of Eq **1.64** can be written as:

$$\sigma = \frac{\hat{\sigma}\pi D^2}{4} \tag{4.4}$$

where $\hat{\sigma}$ is the target radar albedo and D is the target diameter. Assuming that the frequency resolution equals the intrinsic Doppler broadening for an equatorial view of the target, the bandwidth B can be expressed as:

$$B = \frac{4\pi D\cos\delta}{\lambda P_s} \tag{4.5}$$

Table 4.3: Key parameters for simulated sensors.

Tx Band (Freq.) (GHz)	**Antenna Diameter (m)**	**T$_{\text{AMW}}$ (K)**	**Gain (dBi)**
L (1.333)	35	26.68	53.08
S (2.38)	35	18.27	57.90
X (8.56)	35	11.46	68.00
Ka (34.0)	35	19.57	79.77
L (1.333)	70	26.68	59.10
S (2.38)	70	18.27	63.92
X (8.56)	70	11.46	74.02
Ka (34.0)	70	19.57	85.79

where P_s is the rotation period in seconds and δ is the subradar latitude, defined as the angle between the observer line-of-sight and the asteroid's equatorial plane. Here, it is assumed zero (equatorial view) when unknown.

Finally, the bistatic radar equation can be rewritten as:

$$\mathrm{SNR} = \frac{P_R}{\Delta P_N} = \frac{P_T G_T G_R \lambda^2 \hat{\sigma} \pi D^2 \sqrt{\Delta t}\sqrt{\lambda}\sqrt{P_s}}{64\,\pi^3\,4\,R_T^2 R_R^2 k T_{\text{sys}}\sqrt{4\pi D}} = \frac{P_T G_T G_R \lambda^{5/2}}{512\pi^{5/2} k T_{sys}} \cdot \frac{\hat{\sigma} D^{3/2} P_s^{1/2}}{(R_T R_R)^2} \cdot (\Delta t)^{1/2} \tag{4.6}$$

where P_T is the transmitted power, G_T and G_R are the gains of the transmitting and receiving antennas respectively, λ is the signal wavelength, and R_T and R_R are the transmitter-to-target and target-to-receiver ranges respectively. Here the SNR is defined as the ratio of the received signal power to the noise standard deviation after coherent integration over time Δt, rather than to the instantaneous noise power, consistently with the spectral analysis framework described in Section **4.4.2**.

The target population was derived from all NEOs with a close approach to Earth during the solar year 2020 at a distance ≤ 0.2 AU and within the visibility window (elevation$\geq 20°$) of a radar located at the Madrid DSN site. Three size categories were considered: all objects, those with $D \geq 25$ m, and those with $D \geq 140$ m. When some physical properties were unknown, standard values were adopted: radar albedo $\sigma = 0.1$ [193], rotation period $P = 2.1$ h for $D \geq 140$ m and $P = 0.5$ h for smaller objects, and diameter D derived from absolute magnitude H adopting optical albedo 0.14 following the ESA NEOCC database assumptions [194]. A detection threshold of SNR≥ 30/track was adopted. Above this threshold, the detection probability approaches 100% and accurate astrometric measurements can be extracted [186]. A readiness threshold was additionally applied to reflect realistic operational constraints: targets discovered fewer than five days before their visibility window were excluded from the detectable population, since a real-world observing campaign requires planning time for antenna scheduling, frequency coordination, and Doppler prediction preparation.

Figure **4.2** shows the expected yearly detections of NEOs larger than 25 m under average clear weather and very cloudy conditions, for 70 m and 35 m antennas across the four frequency bands, as a function of transmitted power from 0 to 1000 kW. The results reveal several important trends.

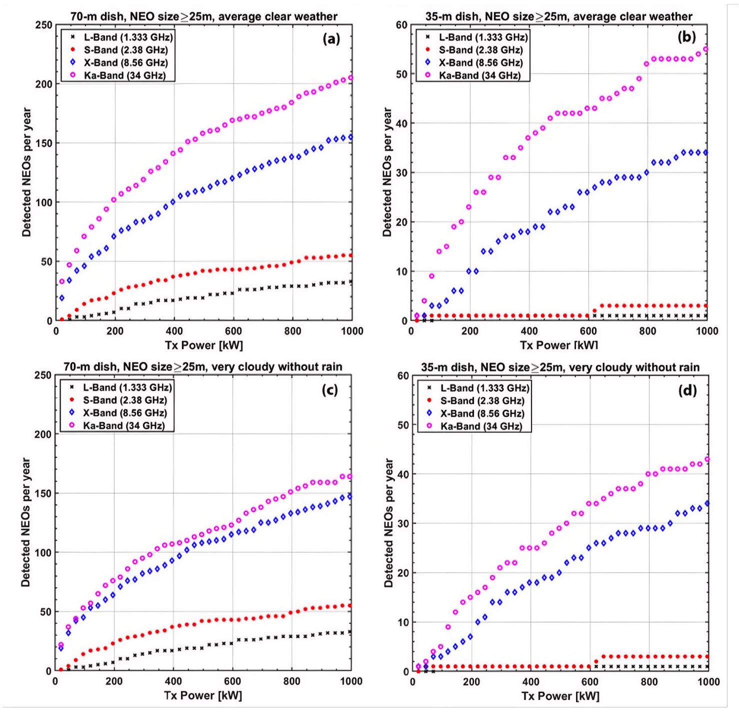

Figure 4.2: Number of expected yearly detections of NEOs greater than 25 m in size by a monostatic radar system as a function of frequency and power if employing a (**a**) 70 m antenna in average clear weather, (**b**) 35 m antenna in average clear weather, (**c**) 70 m antenna in very cloudy weather without rain, and (**d**) 35 m antenna in very cloudy weather without rain. From [155].

The Ka-band consistently outperforms all lower frequencies for a given transmitted power on both antenna sizes and under both weather conditions. This advantage arises from the greater antenna directivity at shorter wavelengths, which means that higher frequencies are intrinsically more efficient for a given aperture and transmitted power. A Ka-band system requires roughly half the transmitted power of an X-band system to achieve the same detection yield. This advantage persists even under cloudy conditions, because the greater sensitivity gain at Ka outweighs the increased atmospheric attenuation. The L- and S-band options offer the practical advantage of being essentially immune to weather variability, but their lower detection rates (particularly for 35-m class antennas) make them significantly less productive. The X-band represents an intermediate choice: it offers substantially better performance than L or S band while benefiting from more mature high-power transmitter technology than Ka-band. Currently, high-power Ka-band systems suffer from significant losses in the amplifiers and in the transmission lines (e.g., in the waveguides) at the power levels required for planetary radar, and these losses were not included in the simulations. The gap between Ka- and X-band performance in practice may therefore be narrower than the ideal calculations suggest.

An additional practical constraint for Ka-band bistatic operations is antenna pointing accuracy. At Ka-band frequencies, the antenna beamwidth narrows significantly, imposing stringent pointing requirements to maintain the target within the main lobe throughout the observation. In general, large radio telescopes require increasingly frequent pointing calibrations at higher frequencies: while at X-band a single calibration at session start is often sufficient, at higher frequencies recalibrations may be needed every $1-2$ hours [195]. In a bistatic planetary radar campaign, where continuous uninterrupted tracking of a moving target is required, such operational constraints further reduce the practical advantage of Ka-band over X-band.

The comparison of the detection rates between 70 m and 35 m antennas illustrates the strong sensitivity advantage of larger apertures, underlining the importance of maximizing aperture if a monostatic European transmitter is to be competitive. However, the 35 m class is more economically realistic for a new-build facility, and the simulation results show that even a 35 m antenna at X or Ka band with several hundred kilowatts of transmitted power could detect tens to over a hundred NEOs larger than 25 m per year under typical weather conditions - a scientifically meaningful contribution.

It is also important to note that bistatic and multistatic configurations were the natural choice for the test campaigns described here for two reasons. First, with the exception of Goldstone's GSSR, DSN antennas are not equipped for monostatic radar observations: they lack the transmit/receive switching hardware and high-power radar-dedicated transmitters required to operate as self-contained radar systems. Second, from a scientific perspective, a primary objective of these campaigns was to assess the contribution of European radio astronomy antennas as receiving elements in a prospective European radar network, evaluating their sensitivity and operational compatibility in a realistic bistatic or multistatic geometry.

4.3.2 European Facilities: The Receiving Network

For the test observation campaigns, four European radio telescopes were selected as receiving stations (Figure **4.3**) on the basis of receiver availability at the transmitted frequencies, scheduling feasibility, absence of strong RFI in the relevant bands, and and instrumental compatibility. In particular, all four antennas are elements of the European VLBI Network (EVN) and share the same Digital Baseband Converter (DBBC) backend for data acquisition and the VLBI Data Interchange Format (VDIF) as common output, together with atomic clocks for synchronization [155]. This standardization simplified the observation procedure and ensured uniform data formats across all receiving stations. Subsequent observational campaigns, including the observation of 2005 LW3 carried out after the project's conclusion, have involved additional receiving stations beyond those. Some of these facilities, that are Effelsberg, the Sardinia Radio Telescope (SRT/SDSA), Medicina, and Noto radio telescopes, have already been introduced in Section **4.2.2**; here we describe their role specifically in the context of the NEO radar network and the operational characteristics that made them suitable for these experiments.

The Effelsberg 100 m telescope (Max-Planck-Institut für Radioastronomie, Bad Münstereifel-Effelsberg, Germany) is the largest fully steerable single-dish radio telescope in Europe. Its 100 m aperture provides a collecting area approximately ten times that of a 32 m antenna, translating directly into superior sensitivity for weak radar echoes from small or distant targets. Effelsberg is equipped with receivers from 300 MHz to 90 GHz, covering all four bands considered in the performance analysis. Its location in a natural valley provides good shielding from RFI, an important consideration for the detection of faint echoes in the presence of anthropogenic radio interferences. However, this local topography imposes azimuth-dependent elevation constraints that must be considered in the observation planning. In the campaigns described below, the high sensitivity of Effelsberg proved critical for obtaining high-resolution spectra with high echo SNR at short integration times during the observation of 2005 LW3. This capability was fundamental for the

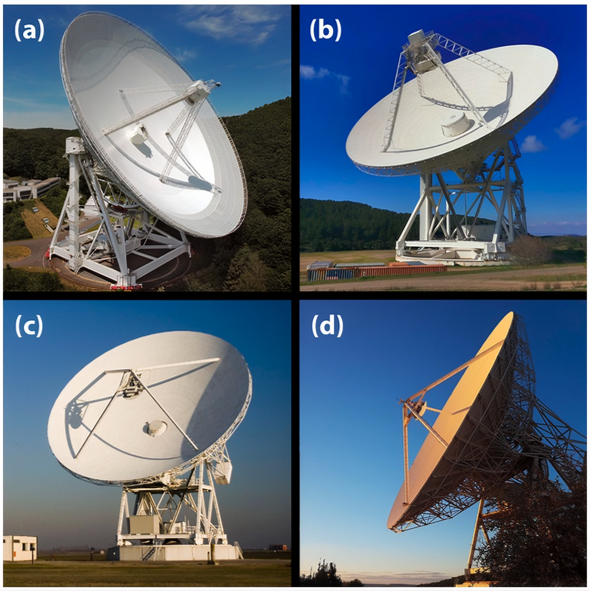

Figure 4.3: The radio telescopes used as receivers in the NEO radar experiments: (**a**) Effelsberg, (**b**) SRT/SDSA, (**c**) Medicina, and (**d**) Noto. From [155].

subsequent 3D shape reconstruction of the target.

The Sardinia Radio Telescope (SRT), operated jointly by INAF and the Italian Space Agency (ASI) is a 64 m active-surface antenna located in the southern Sardinian highlands. The SRT is equipped with a dedicated X-band receiver for deep space and Near-Earth activities [196], making it particularly well-matched to the X-band transmission frequencies employed in all DSS-14 campaigns. Its active surface compensates for gravitational deformation as the antenna moves in elevation, maintaining high aperture efficiency across the operational elevation range.

The Medicina "G. Grueff" Telescope (32 m, INAF, Bologna) and the Noto Telescope (32 m, INAF, Sicily), while significantly smaller than Effelsberg and SRT, are highly versatile instruments with receivers spanning from L to K band. Their role in the network is complementary: they provide additional baselines for multistatic configurations, enabling multi-station SNR combination and cross-validation of detections. Noto's southern location (latitude 36.9° N) provides a slightly lower elevation floor for southern targets compared to the other stations. The Medicina radio telescope has recently been upgraded with an active surface, which improves its efficiency and gain stability at higher frequencies. It should be noted that this upgrade was not yet in place at the time of the observations presented in this work.

The main features of all four receiving antennas are summarized in Table **4.4**.

A fundamental operational advantage of these facilities is that they belong to the European Very Long Baseline Interferometry Network (EVN). This membership ensures that all stations are equipped with hydrogen maser atomic clocks for time and frequency synchronization, and that they employ the same Digital Baseband Converter (DBBC) backend for data acquisition [197]. The DBBC digitizes the downconverted baseband signal and formats it in the VLBI Data Interchange Format (VDIF) - the common raw-data format for all VLBI antennas worldwide [198]. As a consequence, the same post-processing software pipeline can be applied to data from any station in the network without format conversion or station-specific adjustments.

For the transmitting side, the campaigns relied on DSS-14 at Goldstone operating in GSSR configuration (450 kW, X-band) and DSS-63 at Madrid (20 kW, C-high), described in Section **4.2.2**.

Table 4.4: Main features of the receiving antennas. *The Medicina active surface upgrade was completed after the observations presented in this work. †Minimum elevation at Effelsberg is azimuth-dependent, ranging from 5° to over 25° due to the surrounding terrain profile.

	Effelsberg	**SRT/SDSA**	**Medicina**	**Noto**
Location				
Longitude	06°53′01.0″ E	09°14′42.5″ E	11°38′49.0″ E	14°59′20.6″ E
Latitude	50°31′29.4″ N	39°29′30.0″ N	44°31′13.8″ N	36°52′33.8″ N
Antenna				
Diameter	100 m	64 m	32 m	32 m
Rx bands	L, C, X, Ku, K, Ka	P, L, C, X, K	L, S, C, X, K	L, S, C, X, K
Active mirror	yes	yes	yes*	yes
Pointing				
Declination range	−31.4° to +90°	−44.5° to +90°	−40.5° to +90°	−48.1° to +90°
Min. elevation	5°–25°†	6°	5°	5°
Max. elevation	89°	82°	88°	88°
Max. speed Az.	0.50°/s	0.85°/s	0.80°/s	0.80°/s
Max. speed El.	0.27°/s	0.50°/s	0.50°/s	0.50°/s
Pointing acc. (rms)	0.003°	0.001°	0.002°	0.002°

The use of DSS-63 as a transmitter for the 2005 LW3 experiment is noteworthy: it was the first time this antenna had been employed for NEO radar transmission, and it constituted the first fully European-longitude NEO radar campaign after the Evpatoria facility became unavailable. This experiment therefore has not only scientific but also historical significance as a proof of concept for autonomous European planetary radar operations, even if the transmitter power available at DSS-63 (20 kW) is far below what a dedicated European planetary radar would require. The transmitter characteristics are summarized in Table **4.5**. It is worth noting that the physical parameters required for target characterization, such as rotation period and spin axis orientation, are currently derived from observations conducted with Goldstone and Arecibo facilities. While European optical telescopes could in principle contribute to the estimation of some of these quantities, a full radar-based characterization, including delay-Doppler imaging, which requires waveform modulation capabilities not available in the current European bistatic CW configuration, would require a dedicated European high-power transmitter, which does not exist at the time of writing. The European community therefore remains reliant on JPL transmission assets for such measurements.

4.3.3 Observational Results: 2021 AF8 and (4660) Nereus

Below we describe in detail two of the three campaigns that produced the most scientifically relevant results. The 2005 LW3 observation campaign will be introduced in the next section, when we illustrate the progressive development of the observational and data analysis methodology that culminates in the shape reconstruction work of Section **4.5**.

The planning of the NEO radar observations, whose overall workflow is illustrated in Figure

Table 4.5: Main features of the transmitting antennas.

	DSS-14 (Goldstone)	**DSS-63 (Madrid)**
Location		
Longitude	243°06′37.8″ E	355°45′07.2″ E
Latitude	35°25′33.2″ N	40°25′52.4″ N
Antenna		
Diameter	70 m	70 m
Tx bands	X	S, X
Tx power	450 kW	20 kW
Pointing		
Declination range	−34.6° to +90°	−29.6° to +90°
Min. elevation	20°	20°
Max. elevation	89°	89°
Max. speed	0.25°/s	0.25°/s
Pointing accuracy (rms)	0.005°	0.005°

In planetary radar configuration (GSSR).

4.4, involved iterating between two parallel searches: identifying suitable targets and selecting appropriate antenna combinations. On the target side, the starting point was the list of known NEO close approaches from the NEODyS database, filtered for objects with close approach distances ≤ 0.2 AU during the observation period and with predicted SNRs, computed using the simulation tools developed in the first project phase, above the detection threshold for at least one Tx–Rx combination. On the facility side, frequency compatibility, scheduling availability, and absence of significant RFI at the receiving stations had to be verified. The intersection of these constraints produced a list of feasible experiments; the final campaign schedule also accounted for the availability of the transmitting antennas, which were primarily allocated to their deep space communication missions. The complete list of executed experiments since the beginning of the project is given in Table **4.6**. Out of ten attempted observations, seven resulted in detections, a success rate that reflects both the quality of the SNR predictions and the maturity of the post-processing pipeline developed during the project. In the three non-detections (2008 AG33, 1989 JA, and 2015 RN35), the failures were attributed to operational and technical contingencies, including tracking inaccuracies at both the receiving and the transmitting stations, as well as incorrect backend configurations.

Asteroid 2021 AF8 was discovered on 14 January 2021 by the Mt. Lemmon Survey, part of the Catalina Sky Survey program [199]. It is classified as a PHA with an estimated diameter of approximately 300 m derived from its absolute magnitude $H = 20.2$, assuming a mean optical albedo of 0.14. Prior to radar observations, neither its rotation period nor any physical property had been directly measured. The close approach on 4 May 2021 at 0.022 AU provided a good opportunity with a predicted SNR well above the threshold for DSS-14 transmission. The observation was carried out on 3 May 2021 in a bistatic configuration with DSS-14 as transmitter (X-band, 450 kW, CW waveform, RCP polarization) and SRT/SDSA as receiver (LCP polarization). Since DSS-14 was operating in its standard monostatic radar mode the transmitted frequency was

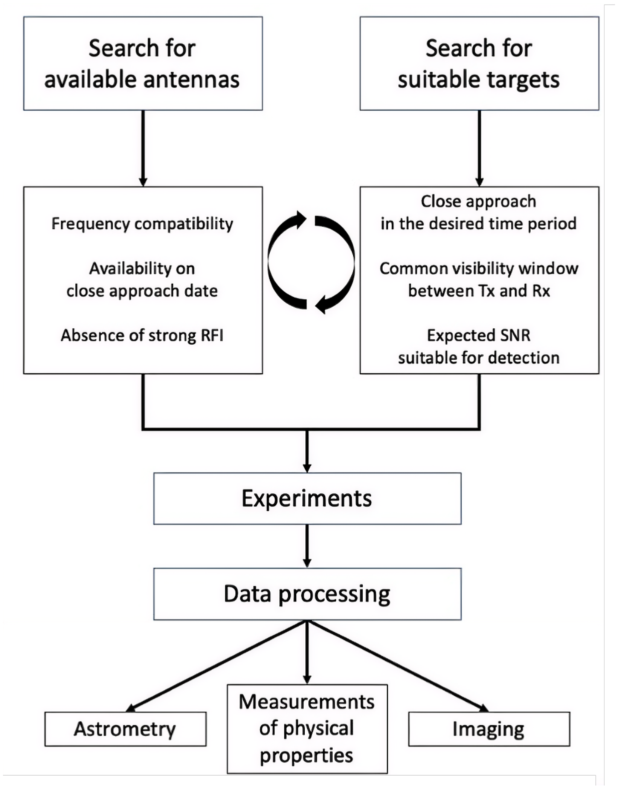

Figure 4.4: Diagram illustrating the phases of the overall activity. From [155].

Table 4.6: List of executed experiments. MC = Medicina, NT = Noto, SRT/SDSA = Sardinia Radio Telescope in Sardinia Deep Space Antenna configuration, EF = Effelsberg, JB = Jodrell Bank Lovell Telescope, EVN = European VLBI Network, e-MERLIN = enhanced Multi Element Remotely Linked Interferometer Network. (∗) VLBI campaign; see Section **4.7**

Date	**Target**	**Tx**	**Rx**	**Detection**
3 May 2021	2021 AF8	DSS-14	SRT/SDSA, MC	Yes (SDSA)
23 Aug 2021	2016 AJ193	DSS-14	MC	Yes
10, 15 Dec 2021	(4660) Nereus	DSS-14	MC	Yes
28 Apr 2022	2008 AG33	DSS-63	EF	No
25 May 2022	1989 JA	DSS-63	EF	No
23 Nov 2022	2005 LW3	DSS-63	MC, EF	Yes
15 Dec 2022	2015 RN35	DSS-63	MC, NT, EF	No
27 Dec 2022	2010 XC15	DSS-63	MC	Yes
25–26 Nov 2024	2006 WB	DSS-14, DSS-63	SRT/SDSA, JB	Yes
18 Sep 2025	2025 FA22 (∗)	DSS-63	EVN + e-MERLIN	Yes

Doppler-compensated for the motion of the target as seen from DSS-14. A signal received at any other location would therefore exhibit a residual Doppler drift due to the geometric difference between the transmitter and receiver positions. Furthermore, because DSS-14 was alternating between transmission and reception in monostatic mode, each cycle of approximately 28 s (close to the round-trip light time to the target) was divided into a transmission phase followed by a reception phase. The spectrograms obtained from the raw SRT data therefore showed both the intermittent signal pattern due to these on/off transmission cycles and the residual frequency drift due to the bistatic geometry.

The Doppler-corrected data enabled the production of high-resolution (0.1 Hz) integrated power spectra (Figure **4.5**), obtained by incoherent integration across the five time intervals containing the signal, for a total integration time of approximately 100 s. The echo was well resolved, enabling measurement of the Doppler broadening via the limb-to-limb width at zero-sigma crossing. Inverting the Doppler bandwidth Equation **4.5**, and assuming an equatorial view ($\delta = 0°$), the measured bandwidth yielded a rotation period estimate of approximately 7.7 h, which under the equatorial assumption represents an upper limit (for $\delta > 0°$ the true period is shorter). Since the SRT receiver acquired only one circular polarization at a time during this experiment, no polarimetric analysis was possible.

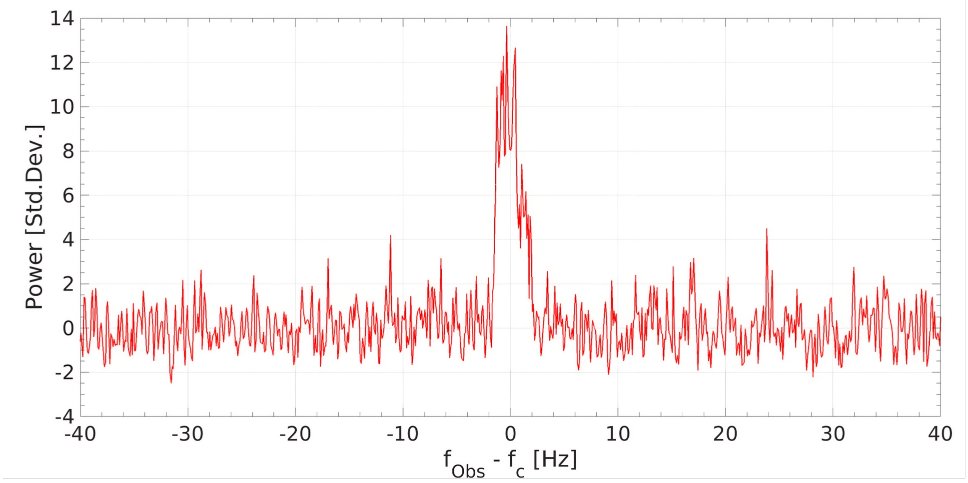

Figure 4.5: High-resolution (0.1 Hz) integrated echo power spectrum of 2021 AF8. Echo power, in units of noise standard deviations, is plotted versus the Doppler frequency (Hz) relative to that of the echo spectral centroid. From [155].

Asteroid (4660) Nereus, discovered in February 1982 [200], is a well-characterized E-class Apollo asteroid with an elongated shape (principal axis dimensions 510 m × 330 m × 240 m), a precisely measured rotation period of $P = 15.16 \pm 0.04$ h, and a known pole direction at ecliptic coordinates $\lambda = +25°$, $\beta = +80°$ (with ±10° uncertainty), all derived from delay-Doppler observations at Goldstone and Arecibo during the 2002 close approach [201]. Its radar albedo of 0.44 and optical albedo of 0.39 are among the highest measured for any asteroid, characteristic of the E spectral class. The December 2021 close approach at 0.026 AU made it an ideal target for validating the European network's capability to reproduce known physical parameters.

During 10–15 December 2021, Goldstone DSS-14 scheduled Nereus observations in Speckle interferometry mode, in which the Very Long Baseline Array (VLBA) was the primary receiving network. The Medicina antenna joined the experiment in eavesdropping mode, acquiring data simultaneously in both LCP and RCP polarizations, in the following time windows:

- 10 December 2021, 12:44:25-13:05:15 UT (approximately 20 minutes);
- 15 December 2021, 12:20:00-12:40:00 UT (approximately 20 minutes)

High-resolution (0.1 Hz) integrated power spectra were produced for both days by combining LCP and RCP (Figures **4.6a** and **4.6b**). To extract the echo bandwidth, both the standard limb-to-limb method (width at zero-sigma crossing) and a multi-parametric model fitting approach [202] were employed. The latter fits simultaneously a linear noise baseline and a signal profile model parametrized by echo width, amplitude, peak position, and shape parameter, via least-squares minimization. Both methods yielded consistent results:

- 10 December 2021: B $= 1.7 \pm 0.1$ Hz;
- 15 December 2021: B $=2.3 \pm 0.2$ Hz.

The subradar latitude on each observation date was computed from the known pole direction: $\delta = +47°(+3°/-7°)$ on 10 December and $\delta = +29°(+9°/-10°)$ on 15 December. Inverting the bandwidth equation with the known rotation period and subradar latitude, the projected polar silhouette breadth was derived:

- 10 December: D $= 379 \pm 54$ m;
- 15 December: D $= 400 \pm 73$ m.

These values are fully consistent with the previously cited principal axis dimensions, confirming that the Medicina measurements correctly recover the known geometry of a well-characterized target.

The detailed description of the signal processing pipeline is given in Section **4.4.1**, where it is developed in full in the context of the 2005 LW3 dataset.

4.3.4 Campaign Results and Scientific Motivation

The test campaigns carried out between 2021 and 2022 demonstrated that European radio telescopes, despite being employed exclusively as receivers and for limited fractions of their scheduled time, are capable of producing scientifically meaningful results from bistatic and multistatic planetary radar observations. All three targets analyzed in this work were successfully detected and characterized at high spectral resolution, validating both the observation planning methodology and the signal processing pipeline developed within the project.

Several technical conclusions emerge from the campaign results. First, the consistency between measured and predicted center-of-mass frequencies at the level of the spectral resolution confirms that the orbital solutions and the Doppler compensation algorithms are reliable. Second, the multiparametric echo profile fitting for bandwidth estimation produced measurements consistent with independently known physical parameters for Nereus, providing a quantitative validation of the data reduction methodology.

At the same time, the campaigns showed the structural limitations of the current European configuration. The absence of a dedicated European transmitter constrains all observations to windows when DSS-14 or DSS-63 are available and willing to collaborate, reducing both flexibility and total accessible observation time. The lack of ranging modulation capabilities at DSS-63 prevents the application of the delay-Doppler techniques, which would provide the range information essential

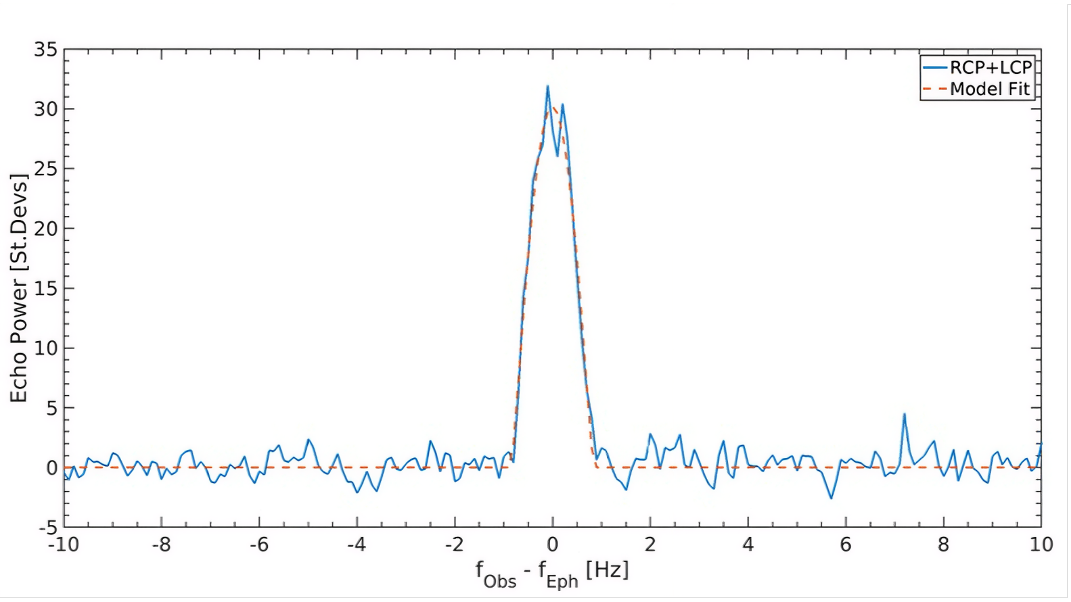

(a) Integrated power spectrum at 0.1 Hz frequency resolution of the Nereus radar echo recorded at Medicina on 2021 December 10 (blue curve). Echo power is plotted in standard deviations of the background noise versus the estimated frequency of the echo from the asteroid's COM. The echo model fit is superimposed on the spectrum (dashed red curve). From [155].

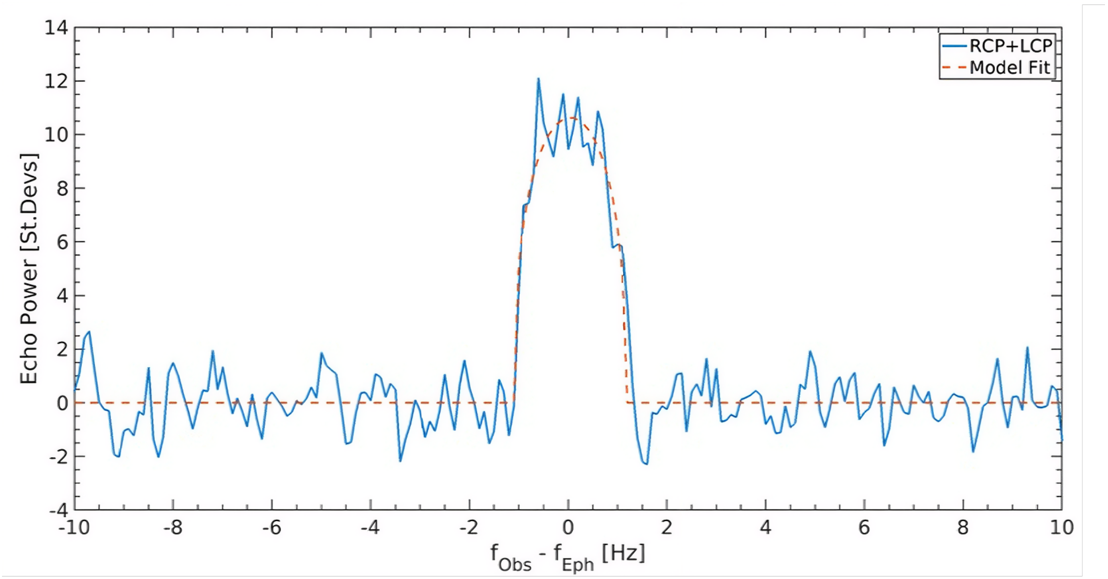

(b) Integrated power spectrum at 0.1 Hz frequency resolution of the Nereus radar echo recorded at Medicina on 2021 December 15 (blue curve). Echo power is plotted in standard deviations of the background noise versus the estimated frequency of the echo from the asteroid's COM. The echo model fit is superimposed on the spectrum (dashed red curve). From [155].

for high-fidelity 3D shape reconstruction.

This last point directly motivates the methodological development presented in the remainder of this chapter. Without range information, shape reconstruction from CW data alone requires algorithms capable of extracting maximum geometric information from the Doppler spectrum and its time evolution. The 2005 LW3 dataset, which will be extensively presented in the following sections, with its high SNR at Effelsberg, long observation arc covering 88% of the asteroid rotation period, and binary nature that introduces an additional detectable component represents the most challenging available test case for such algorithms. The progressive shape reconstruction approach developed in Sections **4.5**, from simple 2D convex projections to adaptive 3D vertex perturbation models, is specifically designed to address this challenge, and to establish a methodology that will remain relevant for future European bistatic campaigns until a dedicated transmitter makes delay-Doppler observations accessible.

4.4 Case study: 2005 LW3

The reconstruction of three-dimensional asteroid shape from CW radar data - Doppler spectra alone, without range information - represents a significantly more ill-posed inverse problem than delay-Doppler inversion. The fundamental challenge is that infinitely many three-dimensional shapes can produce identical Doppler spectra if appropriately oriented and scaled.

The profile of the radar echo Doppler spectrum depends fundamentally on the distribution of line-of-sight velocities across the visible surface, which is determined by the projected shape in the plane perpendicular to the line-of-sight and the rotation velocity field. Consequently, different shapes with identical projections and velocity distributions (up to scaling) can produce identical spectra. Without range information to break this degeneracy, single-epoch CW observations severely underconstrain shape.

Nevertheless, CW observations at multiple epochs spanning an overall observation time comparable to the asteroid rotation period can, in principle, provide sufficient constraints for shape reconstruction. Each epoch samples a different projection of the three-dimensional shape, and the requirement that a single three-dimensional model must simultaneously reproduce all observed Doppler spectra progressively constrains the solution space. The mathematical framework for this inversion, based on progressive refinement from convex hull through perturbative shape models to adaptive vertex representations, is developed and applied to the asteroid 2005 LW3 in Section **4.5**.

Historically, CW radar has been considered unsuitable for shape reconstruction except in highly favorable cases (e.g., slowly rotating objects where multiple viewing geometries can be obtained, or objects with distinctive spectral features). The demonstration that systematic shape reconstruction is achievable from CW Doppler spectra using sophisticated inversion methods represents a meaningful extension of ground-based radar capabilities, particularly valuable in the post-Arecibo era where access to high-power delay-Doppler facilities is increasingly limited.

In this Section, we first present the observation of 2005 LW3 (Section **4.4.1**) described in [155], followed by the data processing techniques (Section **4.4.2**) needed to prepare the dataset for the shape reconstruction algorithms detailed in Section **4.5**.

4.4.1 CW radar observation of 2005 LW3

Asteroid 2005 LW3 was discovered on 5 June 2005 by the Siding Spring Survey [203]. With an absolute magnitude H = 21.6, its diameter was estimated at approximately 170 m before radar observations. The complete list of 2005 LW3 known properties prior to radar observations is reported in Table. **4.7**. The close approach on 23 November 2022 at 0.0076 AU (approximately 1.14 million km) brought it within a favorable radar range, and the predicted SNR was considered sufficient for the DSS-63/Medicina and high for the DSS-63/Effelsberg bistatic radar configurations despite DSS-63's limited 20 kW power. The experiment configuration was fully multistatic: DSS-63 at Madrid transmitted an unmodulated CW signal at 7167 MHz in RCP continuously from 15:50 to approximately 19:00 UT on 23 November 2022 (with brief interruptions for satellite avoidance), while Effelsberg and Medicina recorded simultaneously in linear polarization (LVP and LHP at Effelsberg) and circular polarization (RCP and LCP at Medicina), respectively. The transmitted frequency of 7167 MHz was chosen, in agreement with JPL, to place the received echo in a spectral region free from impacting RFI - a practical constraint that illustrates the careful frequency planning required for real-world campaigns. The system parameters are given in Table **4.8**.

Table 4.7: Orbital and physical properties of 2005 LW3 as known prior to radar observations.

Parameter	**Value**
Orbital parameters	
Epoch (MJD)	60 200.0
Orbit type	Apollo
Eccentricity	0.464
Inclination	6.0°
Perihelion distance	0.772 au
Aphelion distance	2.106 au
Orbital period	630.3 days
Close approach distance	0.0076 au
Close approach date (UT)	23 November 2022, 10:05
Earth MOID	0.00134 au
Physical parameters	
Absolute magnitude H	21.6
Diameter	∼170 m
Rotation period	unknown
Optical albedo	unknown
Radar albedo	unknown
Spectral class	unknown

The raw data recorded at Effelsberg consist of time-series voltage samples in linear polarization (LVP and LHP), digitized at the DBBC backend and stored in VDIF format.

Table 4.8: System parameters for the observation of 2005 LW3.

	DSS-63 (Tx)	**Effelsberg (Rx1)**	**Medicina (Rx2)**
Antenna			
Diameter (m)	70	100	32
Aperture efficiency	0.70	0.55	0.52
Transmitter			
Tx frequency (MHz)	7167	–	–
Tx power (kW)	20	–	–
Tx waveform	CW	–	–
Receiver			
System temperature (K)	–	30–40	90
Polarization	RCP	LVP, LHP	RCP, LCP

The initial stage of the processing pipeline is the Doppler compensation through a phase-stopping procedure [204, 205]. This step is crucial to remove the frequency drift caused by the relative motion between the transmitter, the target, and the receiver. The procedure is structured as follows:

- calculating the received signal frequency as a function of time based on ephemeris, followed by the computation of a polynomial fit that models these frequency variations;
- deriving the phase polynomial fit, the coefficients of which were obtained by integrating the frequency polynomial fit;
- multiplying each raw sample acquired in the time domain by the corresponding complex-valued phase correction derived from the phase polynomial fit.

In the specific case of 2005 LW3 radar observation, the ephemerides were calculated using the JPL orbital solution #30.
To clearly illustrate the impact of the phase-stopping procedure on the received signal, we present a comparison between the spectrograms of the radar echo of 2005 LW3 obtained before and after the application of Doppler compensation. In the uncompensated spectrogram (Figure: **4.7a**), the echo appears as a diagonal streak drifting over time. This tilt corresponds to a drift rate of approximately -7.5Hz/s. In contrast, the spectrogram processed with Doppler compensation (Figure: **4.7b**) shows the echo at a constant frequency and it appears horizontal. Moreover, the compensation significantly afflicts the signal intensity. By examining the color scale values in both spectrograms, it is evident that the echo power in the de-Dopplered data is significantly higher than in the uncompensated case. This is a direct consequence of frequency smearing which afflicts non-compensated data. Frequency smearing not only degrades the SNR but also imposes a fundamental limit on the achievable frequency resolution.

Once the phase-stopping correction has been applied, high-resolution integrated power spectra (Figure **4.8**) are produced by incoherently averaging all spectral frames over the desired integration window (10 minutes), adopting a spectral resolution of 0.1 Hz. Incoherent averaging over N frames

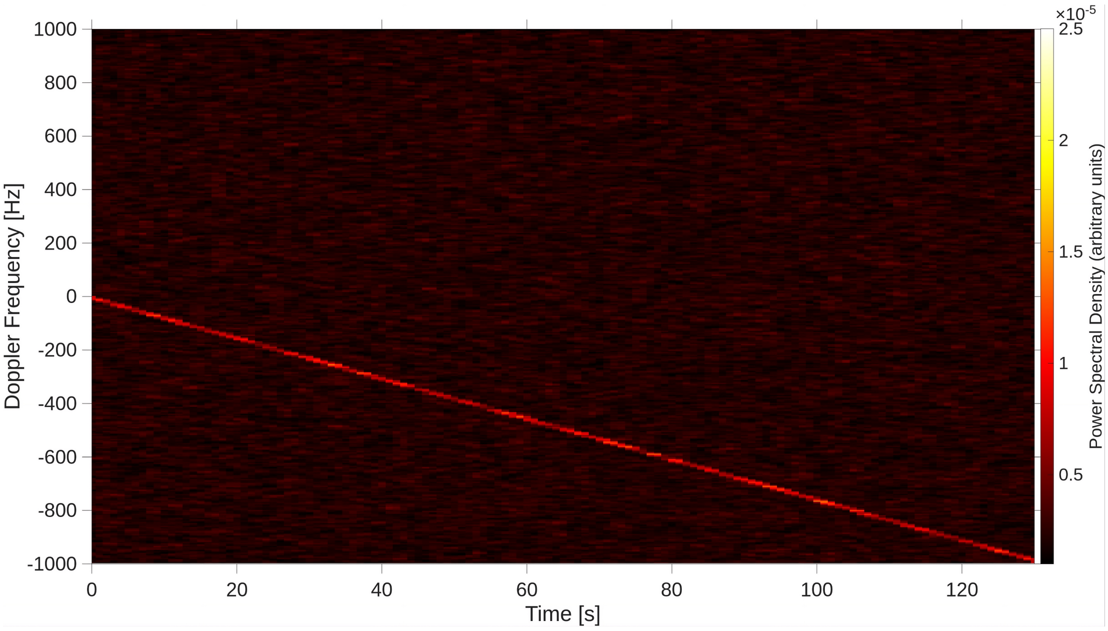

(a) Without Doppler Compensation

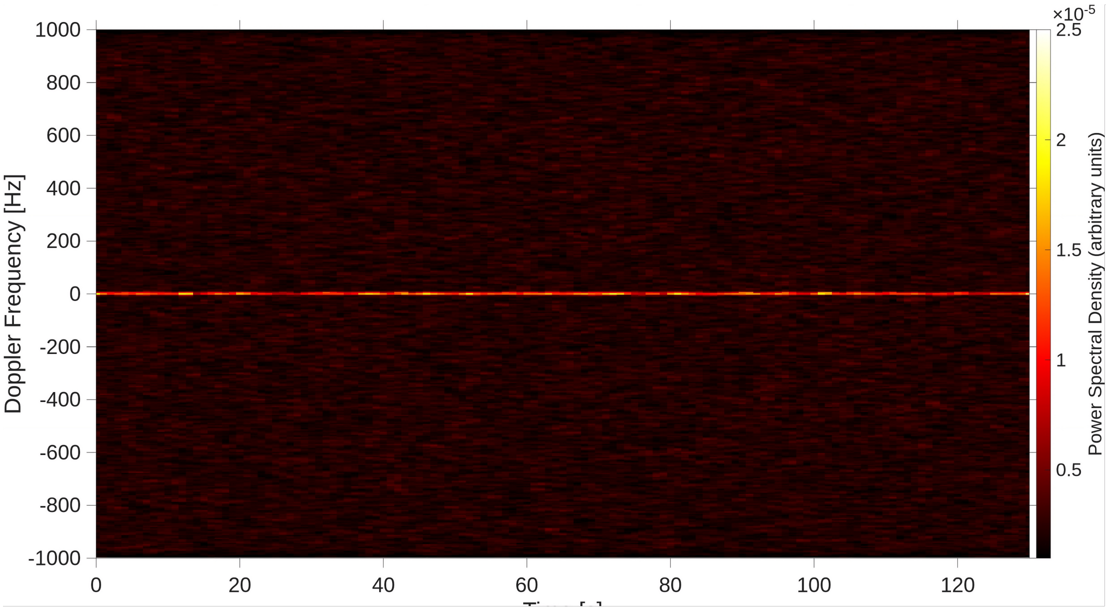

(b) With Doppler Compensation

Figure 4.7: Spectrograms of the signal recorded at Effelsberg for the 2005 LW3 observation, (a) without Doppler compensation and (b) with Doppler compensation. Spectral resolution: 5 Hz, integration time: 2 s. The color scale represents the absolute power spectral density in arbitrary units.

improves the SNR by a factor of $\sqrt{N}$. This integration effectively reduces the speckle noise and allows the characterization of spectral features that would be below the noise threshold in individual, short-integration frames. The echo power is expressed in units of the standard deviation of the background noise, estimated from the off-echo spectral channels. The most striking discovery was the binary nature of the system. In the high-resolution (0.1 Hz) power spectra acquired at Effelsberg, a narrow secondary peak was visible at approximately $+4$ Hz relative to the expected frequency from the center of mass, superimposed on the broad echo of the primary body (Figure **4.8**). The same secondary peak was also detected at Medicina, albeit requiring a longer integration time and a coarser frequency resolution owing to the lower sensitivity of the 32-m antenna relative to Effelsberg. The consistency of the asteroid's satellite detection across two independent receiving stations with different apertures and polarization configurations confirms that the feature is a genuine physical signal rather than an instrumental artifact.

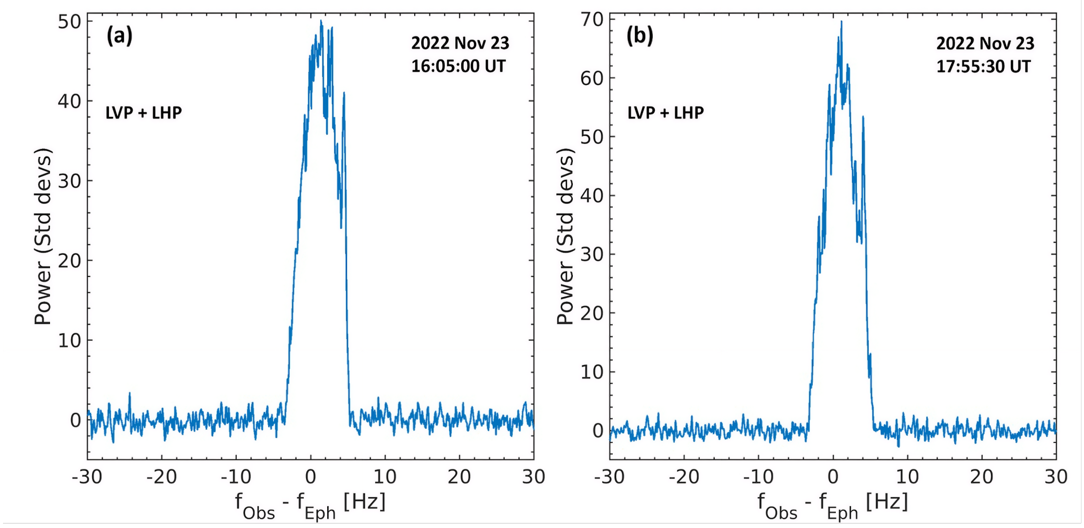

Figure 4.8: Integrated power spectra at 0.1 Hz frequency resolution, 10 min integration time, of 2005 LW3 echo recorded at Effelsberg on 23 November 2022 at (**a**) 16:05:00 UT and (**b**) 18:55:30 UT. Echo power on the vertical axis is plotted in standard deviations of the background noise. The Doppler frequency on the horizontal axis is relative to the ephemeris-based estimate frequency of the echo from the asteroid COM. The spike at ~ 4 Hz in both spectra is the echo from the satellite.

Goldstone delay-Doppler imaging (Figure **4.9**) subsequently confirmed the binary nature and revealed that the primary body is approximately 400 m in size - more than twice the pre-radar diameter estimate - with the satellite at a separation of approximately 4000 m and an equatorial diameter between 50 and 100 m [206].

As additional proof, the full-track integrated power spectrum of 2005 LW3 recorded both at Effelsberg and Medicina, obtained by summing incoherently all power spectra acquired during the entire observation arc (Figure **4.10**), reveals the secondary peak persisting even after averaging over nearly the entire asteroid rotation, confirming the satellite echo detection.

Additionally, all echo profiles showed a center-of-mass frequency displacement of 1.0 ± 0.1 Hz relative to the ephemeris prediction. This offset corresponds to a range-rate residual that provides a meaningful astrometric constraint for orbit refinement: a direct demonstration of the astrometric contribution that European bistatic observations can make even with limited transmitter power.

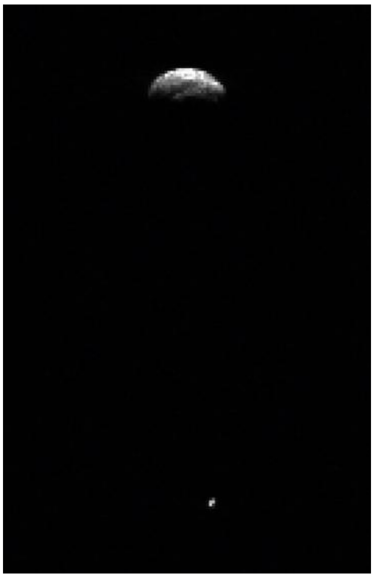

Figure 4.9: Delay-Doppler image of 2005 LW3 and its satellite obtained at Goldstone (NASA/JPL).

As mentioned above, the Medicina radio telescope detected the 2005 LW3 echo in both the same (SC) and opposite (OC) circular polarization relative to the transmitted signal (see Figure **4.11**), so it was possible to measure a low circular-polarization ratio between 0.1 and 0.2. As pointed out in Section **4.1.1**, this ratio is a crucial observable in NEO radar observations for the target physical characterization, as it is related to the NEO surface and sub-surface roughness at the wavelength scale [168]. Specifically, a low polarization ratio like the one estimated for 2005 LW3 indicates that the object appears smooth at the centimeter scale. To improve the accuracy of the polarization ratio measurements, relatively low-resolution spectra (1 Hz frequency resolution) were produced. This approach maximized the SNR and minimized speckle noise by concentrating the signal into wider frequency bins, ensuring a more reliable calculation of the integrated power for each polarization. However, this lower resolution precludes the detection of the satellite echo,

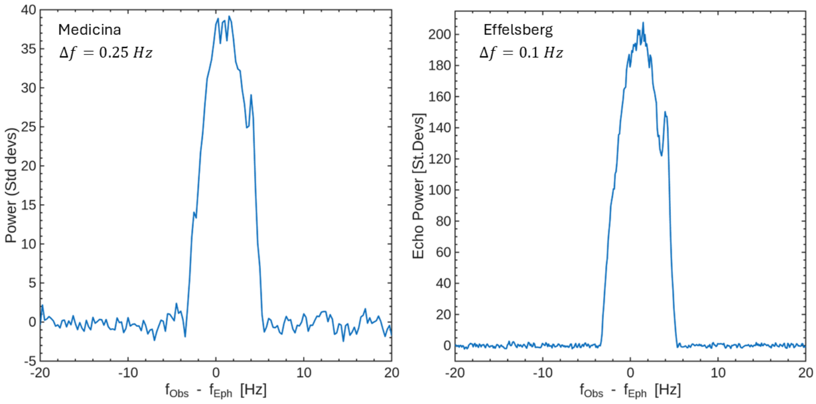

Figure 4.10: Full-track integrated echo power spectra of 2005 LW3. Left: Medicina data (0.25 Hz frequency resolution, OC polarization only). Right: Effelsberg data (0.1 Hz frequency resolution, combined LVP and LHP polarizations). The integration time for each spectrum is $\sim$ 2.5 h. Echo power is plotted in standard deviations of the background noise versus the estimated echo frequency relative to that from the asteroid's COM. From [155].

clearly visible in the high-resolution spectra shown previously.

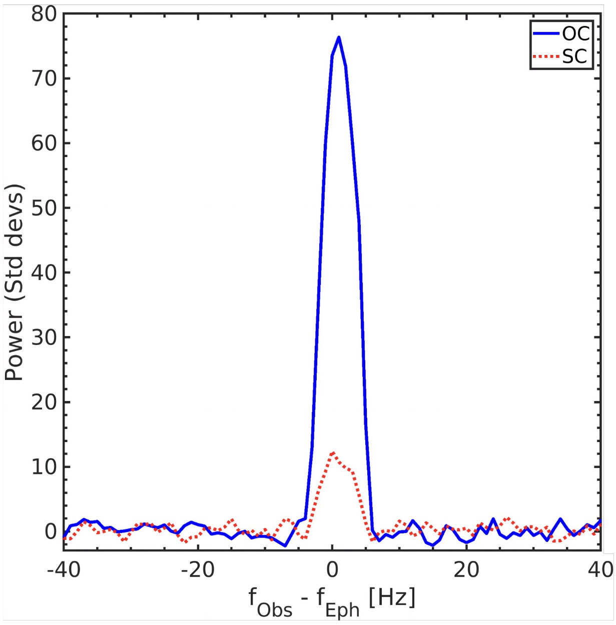

Figure 4.11: Echo power spectrum of 2005 LW3 in the OC (solid blue) and in the SC (dotted red) polarization derived from Medicina data. Frequency resolution: 1 Hz, full-track integration.

The echo, detected from both receiving stations, was well resolved in the frequency domain, and from the limb-to-limb Doppler bandwidth, which defines the bandwidth as the frequency interval between the zero-crossings of the echo profile above the noise baseline, a rotation period of approximately 4.0 h was measured (assuming equatorial view), consistent with independent Goldstone delay-Doppler observations.

4.4.2 Spectral analysis for shape modeling

The data reduction pipeline extracts the edge frequencies f_+ and f_- of the radar echo from each integrated power spectrum, providing a direct measurement of the Doppler bandwidth of the asteroid. These quantities encode the projected rotational velocity of the target's limbs and serve as the primary input for the shape reconstruction algorithms described in Section **4.5.1**. Specifically, the pipeline estimates the values of the support function of the asteroid's polar silhouette and their associated uncertainties. The pipeline was implemented in MATLAB and applied to each spectrum in the 2005 LW3 dataset independently, iterating over the full observation arc. Each input file contains a power spectrum (power spectral density as a function of baseband frequency), together with metadata including the Julian date of the spectrum midpoint, the integration time, and the spectral resolution.

The first processing step consists of estimating and removing the noise baseline from each spectrum. The baseline is modeled with a fifth-degree polynomial fit computed over two spectral regions flanking the echo. These off-echo regions are chosen to be sufficiently far from the echo to be uncontaminated by signal, while remaining close enough to sample the local noise level accurately. The polynomial is evaluated over all spectral channels in the region of interest and subtracted from the spectrum.

After background subtraction, the residual spectrum is normalized by the standard deviation of the off-echo noise, converting the power spectral density into units of noise standard deviations.

The normalization process is crucial for putting all the observed spectra into a common, physically meaningful scale. All the subsequent processing operates on this normalized representation.

As described in Section **4.4.1**, the 2005 LW3 system is binary: the radar echo contains a narrow spectral spike from the satellite superimposed on the broader echo of the primary body. For the purpose of estimating the shape of the primary, the satellite contribution must be removed before any analysis of the echo bandwidth or profile. The satellite echo occupies a narrow and precise frequency interval in each spectrum. These intervals may vary across spectra because the satellite's Doppler frequency relative to the primary's center of mass changes as the mutual orbit evolves over the observation arc. The removal consists of extracting and storing the satellite spectral channels separately, then extracting them from the working copy of the spectrum by replacing it with a parametric fit of the main echo in the satellite's echo regions. Subsequent steps operate on this satellite-free version, which contains the primary echo without contamination from the secondary. The satellite channels are retained for separate analysis if needed.

The model employed to fit the primary echo profile is the one described in [202], and already used in the [155] work. According to the Jurgens model, the echo power spectrum can be represented as a uniformly scattering sphere as:

$$S(f) = S_0 \left[1 - \left(\frac{2(f - f_{\text{mod}})}{B} \right)^2 \right]^n \tag{4.7}$$

for $|f - f_{\text{mod}}| \leq B/2$, and zero otherwise, where S_0 is the peak amplitude, f_{mod} is the frequency of the echo centroid, B is the echo bandwidth at zero-amplitude crossing, and n is a shape parameter controlling the curvature of the spectral profile. Despite the spherical assumption, this model provides a flexible description of the echo profile and it is widely used in planetary radar data reduction analysis as a robust echo shape estimator. The fit is performed using a nonlinear least-squares routine, which minimizes the sum of squared residuals between the observed (satellite-free, normalized) spectrum and the model. The four parameters $(S_0, f_{\text{mod}}, B, n)$ are physically constrained: $S_0 > 0$; f_{mod} is limited to the frequency interval containing the echo; $B > 0$ and bounded above by a multiple of the initial estimate; and $n \in [-1, 10]$, allowing both flat-top and peaked profiles. This range represents the standard bounds adopted in planetary radar spectral fitting [202]; the actual value of n is determined iteratively by the nonlinear least-squares routine and varies from spectrum to spectrum depending on the observed echo shape. Uncertainties on the four fitted parameters are estimated from the Jacobian matrix using the method described in [207], which yields the 95% confidence interval.

The fitted model is evaluated over the baseband spectral region of interest (ROI), which encompasses the echo and the surrounding background noise, including the satellite echo channels. The reason behind this is twofold: it provides a clean representation of the expected primary echo shape for visual comparison with the data, and it defines the model-based estimator of the echo edge frequencies described below.

The geometrically significant quantity that this data reduction pipeline is designed to extract from each spectrum is the pair of frequencies f_+ and f_- corresponding to the echo edges: the Doppler shifts of the approaching and receding limbs of the asteroid. These frequencies are related to the geometry of the asteroid's polar silhouette through the support function, as described in Section **4.5.1**. Two independent estimators of f_+ and f_- are computed.

The first estimator uses the fitted Jurgens model. Given the four parameters $(S_0, f_{\text{mod}}, B, n)$ the edge frequencies at a specified sigma threshold σ_{cross} are found by solving $S(f) = \sigma_{\text{cross}}$ numerically

using Brent's method, starting from the zero-crossing limits of the model as initial guesses. This gives $f_{\text{model},+}$ and $f_{\text{model},-}$, together with the model bandwidth $B_{\text{model}} = |f_{\text{model},+} - f_{\text{model},-}|$.
The second estimator operates directly on the observed normalized spectrum, without model assumptions. For each edge, the two spectral channels straddling the σ_{cross} threshold are identified - one above and one below - and the crossing frequency is estimated by linear interpolation:

$$f_{\pm} = f_1 + \frac{\sigma_{\text{cross}} - S_1}{m} \tag{4.8}$$

where S_1 and f_1 are the power and frequency of the channel just inside the echo profile, S_2 and f_2 are those of the channel just outside, and $m = (S_2 - S_1)/(f_2 - f_1)$ is the local slope. The uncertainty on each crossing frequency is propagated analytically from the spectral resolution Δf, the noise standard deviation $\delta S = 1$ and the local slope m:

$$\delta f_{\pm} = \sqrt{\Delta f^2 + \left(\frac{\delta S}{m}\right)^2 + \left(\frac{|\sigma_{\text{cross}} - S_1|}{m^2} \cdot \frac{\sqrt{2}\sqrt{m^2(\Delta f/s)^2 + \delta S^2}}{|f_2 - f_1|}\right)^2} \tag{4.9}$$

This expression accounts for three independent contributions: the spectral quantization (Δf), the noise on the power measurement ($\delta S/m$), and the propagation of the slope uncertainty through the interpolation formula. In the limiting case of a steep slope (large $|m|$) and fine spectral resolution, the uncertainty reduces to the spectral resolution alone; for shallow slopes or coarse resolution, the noise contribution dominates. The threshold σ_{cross} is a free parameter of the pipeline, set to $\sigma_{\text{cross}} = 2$ standard deviations in the analyses presented here, corresponding to a well-defined trade-off between sensitivity to the true echo edge and noise excursions [208]. Figure **4.12** provides a schematic illustration of the two estimators of f_+ and the relevant quantities involved.
From the edge frequencies $f_{\pm}$, the corresponding values of the support function are computed. The support function $p(\theta)$ of the convex hull of the asteroid's polar silhouette is defined in [209] as the distance from the rotation center to the tangent line to the hull perpendicular to the radar line of sight at rotational phase θ. By the geometric relation between echo bandwidth and target shape, the two edge frequencies f_+ and f_- of each spectrum correspond to the two support function values:

$$p_+ = \frac{\lambda P |f_+|}{4\pi}, \quad p_- = \frac{\lambda P |f_-|}{4\pi} \tag{4.10}$$

being λ the radar wavelength and P is the rotation period. In physical units, $p_{\pm}$ has dimensions of meters multiplied by $\cos\delta$, where δ is the subradar latitude. Hence, the support function values evaluated here represent projections of the hull dimensions onto the equatorial plane, with the cos δ treated as a systematic scale uncertainty. Uncertainties ϵ_p are propagated directly from Equation **4.9** through the same linear relation:

$$\epsilon_{p_{\pm}} = \frac{\lambda P}{4\pi} \cdot \delta f_{\pm} \tag{4.11}$$

Each spectrum is assigned a rotational phase θ based on the Julian date jd_{mid} of its midpoint and a reference epoch jd_0 (the "zero-phase time"). The phase, accounting for its periodicity, is computed as:

$$\theta = \left(360^\circ \cdot \frac{\Delta t}{P}\right) \pmod{360^\circ} \tag{4.12}$$

where $\Delta t = (jd_{\text{mid}} - jd_0) \times 86400$ s is the elapsed time since the reference epoch in seconds, and P the rotation period in seconds. Each spectrum contributes two support function measurements: p_+

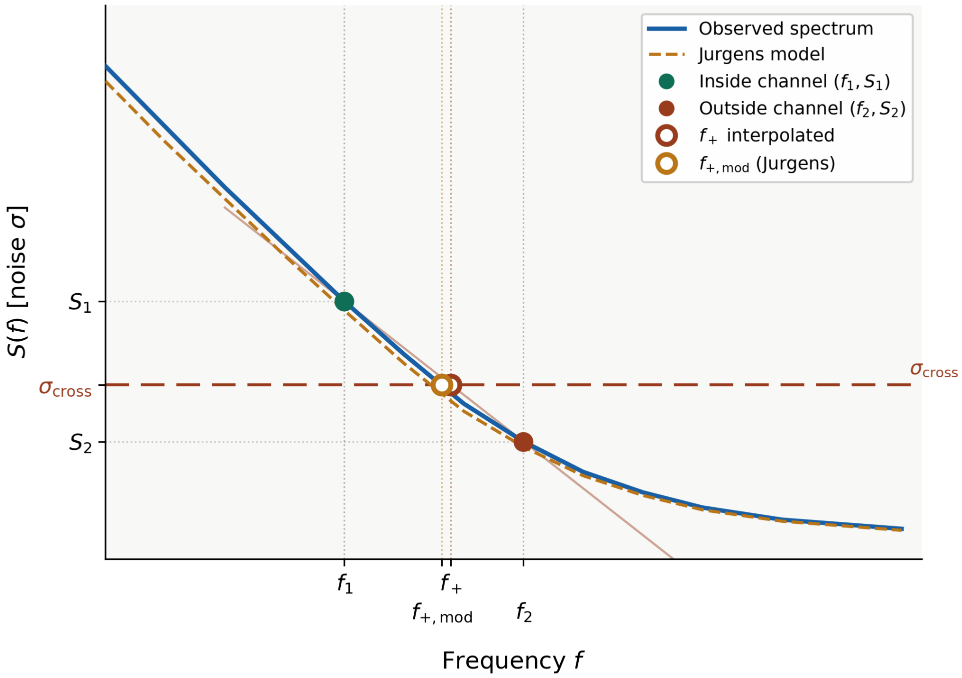

Figure 4.12: Zoom on the right edge of a schematic Doppler echo spectrum, illustrating the two estimators of f_+. The inside channel (f_1, S_1) lies above σ_{cross}, while the outside channel (f_2, S_2) lies below it. The interpolated crossing f_+ (open red circle) is obtained by linear interpolation of the observed spectrum (blue) between the two straddling channels, while $f_{+,\text{mod}}$ (open amber circle) is the crossing of the fitted Jurgens model (dashed amber) with σ_{cross}.

at phase θ (corresponding to the approaching limb at that rotational orientation) and p_- at phase $\theta + 180°$ (corresponding to the receding limb, which is geometrically equivalent to the approaching limb seen from the opposite side). The integration time of each spectrum corresponds to a phase span:

$$\Delta\theta = 360° \cdot \frac{t_{\text{int}}}{P} \tag{4.13}$$

the start and end phases of each integration window are also recorded.

At the end of the pipeline, the measurements from all spectra in the dataset are collected into a set of vectors sorted by rotational phase. Each element of these vectors represents one support function measurement and carries: the support function value p and its uncertainty ϵ_p, both in units of meters × cos δ; the rotational phase θ at the midpoint of the spectrum's integration window (10 minutes integration time); the phase span $\Delta\theta$ and the Julian date of the spectrum midpoint.

Figure **4.13** shows the 2005 LW3 spectrum at the beginning of this pipeline and its resulting one at the end, for the same rotational phase θ. This dataset constitutes a noisy, discretely sampled representation of the support function $p(\theta)$ of the asteroid's polar silhouette, which will be the basis of the 2D convex hull algorithm described in the next section.

4.5 Progressive 3D shape reconstruction algorithms

4.5.1 2D convex hull model

The 3D shape reconstruction pipeline begins with the estimation of the convex hull H of the asteroid's polar silhouette, defined as the two-dimensional convex envelope of the body's projection

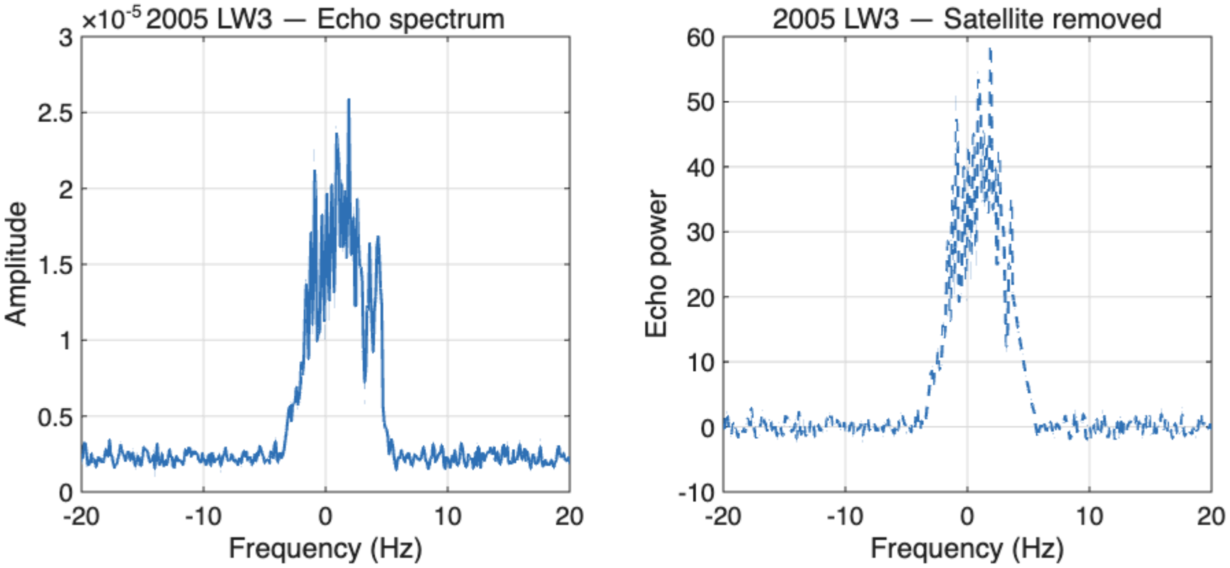

Figure 4.13: **Left**: 10 minutes integrated 2005 LW3 spectrum at 0.1 Hz. **Right**: 10 minutes integrated 2005 LW3 spectrum at 0.1 Hz, after the removal of the satellite's echo. The spectrum has been baseline-subtracted using a fifth-degree polynomial fit to the noise regions and normalized to units of noise standard deviation.

onto the equatorial plane, as seen along the spin vector. The theoretical basis for this reconstruction was established by [209] and the implementation described here followed their approach directly. The geometric consideration that is the basis of this method is that the instantaneous echo power spectrum of a rotating asteroid carries information about the breadth of its polar silhouette measured normal to the radar line-of-sight. This relation is illustrated in Figure **4.14**. Specifically, the Doppler frequencies f_+ and f_- of the approaching and receding echo edges correspond, at each rotational phase angle θ, to the distances $p_+(\theta)$ and $p_-(\theta)$ from the plane containing the spin vector and the line-of-sight to the tangent planes touching the approaching and receding limbs of the body. As shown in Section **4.4.2**, these distances are the values of the support function of the convex hull H:

$$p(\theta) \equiv p_+(\theta) \equiv p_-(\theta + 180^\circ) \tag{4.14}$$

that is a 2π period function of the rotational phase θ. $p_+(\theta)$ measures the extent of the silhouette in the direction θ, while $p_-(\theta)$ measures it in the opposite direction $\theta + 180^\circ$, so that the two edge frequencies together sample the full support function over a complete rotation. According to [210], the support function completely determines the convex hull H: given $p(\theta)$, the Cartesian coordinates of the boundary of H in the body-fixed frame rotating with the asteroid are:

$$x(\theta) = p(\theta)\cos\theta - p'(\theta)\sin\theta \tag{4.15}$$

$$y(\theta) = p(\theta)\sin\theta + p'(\theta)\cos\theta \tag{4.16}$$

where $p'(\theta)$ is the derivative of the support function with respect to θ. The radius of curvature of H at the point where the tangent plane is perpendicular to direction θ is:

$$r(\theta) = p(\theta) + p''(\theta) \tag{4.17}$$

If H is a convex curve, the radius of curvature must be non-negative everywhere: $r(\theta) \geq 0 \ \ \forall\theta$. This is the convexity constraint that distinguishes hull estimation from unconstrained least-squares

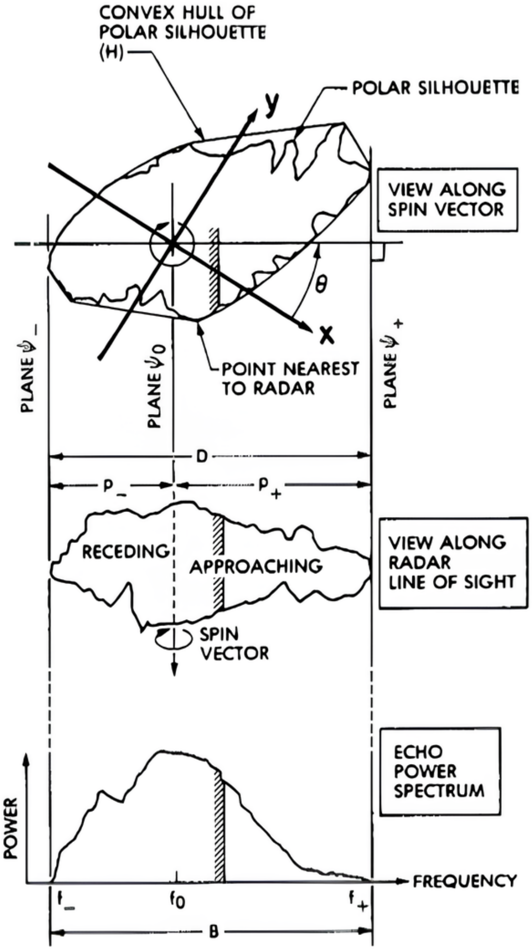

Figure 4.14: Geometric relationships between an asteroid's shape and its echo power spectrum. The plane ψ_0 contains the line-of-sight and the asteroid's spin vector. Echo from any portion of the asteroid intersecting ψ_0 has Doppler frequency f_0. The cross-hatched strip of power in the spectrum corresponds to echoes from the cross-hatched strip on the asteroid. The x, y coordinate system (dark lines) rotates with the asteroid. From [209].

fitting.

The support function is represented as a truncated Fourier series with M harmonics:

$$p(\theta) = a_0 + \sum_{m=1}^{M} (a_m \cos m\theta + b_m \sin m\theta) \tag{4.18}$$

$$p'(\theta) = \sum_{m=1}^{M} m(-a_m \sin m\theta + b_m \cos m\theta) \tag{4.19}$$

We define the parameter vector containing all Fourier coefficients as:

$$\mathbf{x}^T = (a_0, a_1, b_1, a_2, b_2, ..., a_M, b_M) \tag{4.20}$$

of length $2M+1$. For L input spectra, the preprocessing pipeline of Section **4.4.2** yields 2L support function measurements assembled into the data vector $\mathbf{p}^T = (p_1, p_2, ..., p_{2L})$, each measured at a rotational phase θ_i. A necessary condition for a meaningful reconstruction of H is that the 2L support function measurements cover the full rotation of the asteroid, ideally without large gaps in phase. This requirement follows directly from the Fourier nature of the problem: the coefficients a_m and b_m of the support function are determined by the data through a weighted least-squares fit, and any unsampled interval of rotational phase corresponds to a portion of the hull that is entirely unconstrained by the observations. A gap in phase coverage will cause the estimated hull to collapse toward an unphysically small or irregular shape in the unsampled directions. In practice, coverage of most of the full rotation is considered the minimum for a reliable hull estimate. The 2005 LW3 dataset satisfies this requirement: as noted in Section **4.4.1**, the observation arc spans approximately 88% of one full rotation period, providing near-complete phase sampling with no large gaps.

The design matrix $\mathbf{A}$ of size $2L \times (2M+1)$ has its i-th row given by:

$$\mathbf{A}_{i,:} = (1, \cos\theta_i, \sin\theta_i, \cos 2\theta_i), \sin 2\theta_i, ..., \cos M\theta_i, \sin M\theta_i) \tag{4.21}$$

such that the Fourier model evaluated at the observation phases becomes: $\hat{\mathbf{p}} = \mathbf{A}\mathbf{x}$.

The convexity constraint of Equation **4.17** translates, after computing the second derivative of the Fourier series, into a set of linear inequality constraints on $\mathbf{x}$. Specifically, for $r(\theta) \geq 0$ to hold at N discrete phases $\theta_k = 2\pi k/N$, the constraint matrix $\mathbf{F}$ of size $N \times (2M+1)$ has its k-th row given by:

$$\mathbf{F}_{k,:} = (1, 0, 0, (1-2^2)\cos 2\theta_k, (1-2^2)\sin 2\theta_k, ..., (1-M^2)\cos M\theta_k, (1-M^2)\sin M\theta_k) \tag{4.22}$$

The full set of inequality constraints is then $\mathbf{F}\mathbf{x} \geq \mathbf{0}$.

If the measurement errors ϵ in $\mathbf{p}$ are zero-mean random variables with diagonal covariance matrix $\mathbf{V}_p = \mathrm{diag}(\epsilon_{p1}^2, ..., \epsilon_{p2L}^2)$, the optimal estimate of $\mathbf{x}$ is obtained by minimizing the weighted sum of squared residuals:

$$Q(x) = (\mathbf{p} - \mathbf{A}\mathbf{x})^T \mathbf{V}_p^{-1} (\mathbf{p} - \mathbf{A}\mathbf{x}) \tag{4.23}$$

subject to the convexity constraint. Expanding $Q(x)$ and discarding the constant term, this is equivalent to minimizing the quadratic form:

$$\frac{1}{2}\mathbf{x}^T\mathbf{B}\mathbf{x} - \mathbf{b}^T\mathbf{x}, \quad \mathbf{B} = \mathbf{A}^T\mathbf{V}_p^{-1}\mathbf{A}, \quad \mathbf{b} = \mathbf{A}^T\mathbf{V}_p^{-1}\mathbf{p} \tag{4.24}$$

This is strictly a convex quadratic programming problem [211], which has a unique global minimum. If the unconstrained solution $\hat{\mathbf{x}} = \mathbf{B}^{-1}\mathbf{b}$ already satisfies all the inequality constraints, it is also the constrained solution. If not, the constraints push the solution toward the nearest point on the feasible boundary. Geometrically, it means toward the most "ellipsoidal" hull consistent with the data.

In our implementation, the unconstrained least-squares solution is used as the initial guess, and the constrained minimum is found using the sequential least-squares programming (SLSQP) solver. The number of Fourier harmonics is set to $M = 10$, following the original choice of [209], which captures the dominant shape features while limiting the risk of overfitting noisy data. The convexity constraint is enforced at $N = 2L$ phases (coinciding with the observation phases), which is sufficient to guarantee convexity over the entire rotation. $\mathbf{V}_p$ is constructed as a diagonal matrix with the uncertainties ϵ_p on the diagonal, which is equivalent to weighting each measurement inversely proportional to its uncertainty; this choice gives more uniform weighting across measurements with very different noise levels and is consistent with the spectral analysis described in Section **4.4.2**.

Once the optimal coefficient vector $\hat{\mathbf{x}}$ is determined, the hull H is reconstructed by evaluating the parametric Equations **4.15** over a dense grid of phases $\theta \in [0, 2\pi]$. The support function and its derivative are computed analytically from Equations **4.18** that lead to Equations **4.15**. The resulting curve is the pole-on projection of the convex envelope of the asteroid's equatorial silhouette, in units of $m/\cos\delta$, where δ is the subradar latitude as defined in Section **4.3.1**.

From the hull, two equivalent semi-axes are computed:

$$a_{\mathrm{equiv}} = \frac{\max x - \min x}{2}, \quad b_{\mathrm{equiv}} = \frac{\max y - \min y}{2} \tag{4.25}$$

representing the maximum extents of the hull in the two orthogonal directions in the equatorial plane. These quantities are related to the support function values p_+ and p_- through the parametric Equation **4.15**: $x(\theta)$ and $y(\theta)$ are computed directly from $p(\theta)$ and its derivative $p'(\theta)$, with the hull boundary coordinates varying with rotational phase and encoding contributions from p_+ and p_- measurements at all phases. a_{equiv} and b_{equiv} are therefore representative estimates of the asteroid's equatorial dimensions, summarizing the full set of support function measurements. Their ratio a/b provides a model-independent lower bound on the axis ratio of the asteroid's equatorial cross-section, since the hull is the convex envelope and therefore always larger than or equal to the true silhouette in every direction. Figures **4.15a** and **4.15b** schematically illustrates the full reconstruction pipeline, from the input spectra to the reconstructed convex hull, showing how individual phase measurements map to specific points on the hull boundary.

The 2D convex hull model is the most directly constrained by the data and makes the fewest assumptions: it requires only a known rotation period and the measured support function values, without any hypothesis about the three-dimensional shape. However, it carries three fundamental limitations. First, by construction, it yields only the convex envelope of the polar silhouette, with all concavities "filled in" - any non-convex features of the true shape are invisible to this method. Second, it provides no information about the three-dimensional figure of the body: the hull is a purely equatorial projection, and the out-of-plane shape remains unconstrained. Third, as already mentioned during this analysis, the method requires that the observed spectra cover almost the full rotation of the asteroid, since any rotational phase for which no spectrum is available constitutes a blind sector of the hull, introducing a degeneracy that the Fourier fit cannot resolve. These limitations motivate the successive levels of the reconstruction hierarchy described in the following sections.

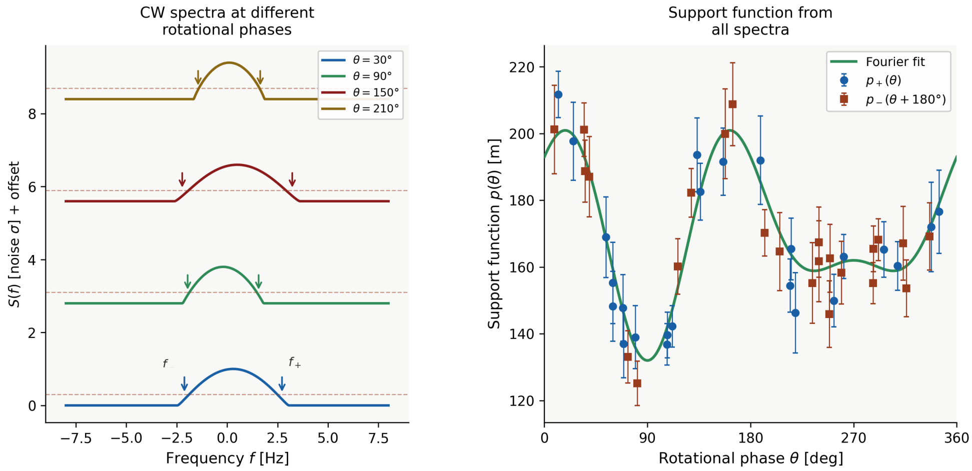

(a) Illustrative example of the data reduction pipeline. **Left**: CW power spectra at four different rotational phases θ, normalized to units of noise standard deviation. For each spectrum, the edge frequencies f_+ and f_- (arrows) are extracted at the σ_{cross} threshold (red dashed line). **Right**: support function values $p_+(\theta)$ (circles) and $p_-(\theta + 180°)$ (squares) derived from all spectra, plotted as a function of rotational phase. The green solid line shows the fitted Fourier series.

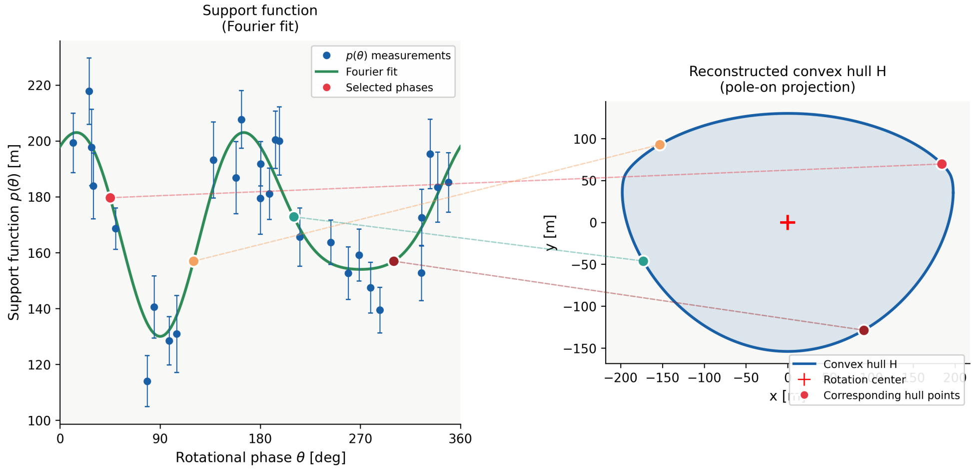

(b) Illustrative example of the convex hull reconstruction pipeline. **Left**: measured support function values $p(\theta)$ (blue circles with error bars) and the fitted Fourier series (green line) as a function of rotational phase θ. Four selected phases are highlighted in color. **Right**: pole-on projection of the reconstructed convex hull H. The colored points on the hull correspond to the same rotational phases highlighted in the left panel, illustrating the geometric mapping from support function values to hull boundary points. The red cross marks the rotation center.

4.5.2 Convex Hull results

The convex hull reconstruction of 2005 LW3, obtained from the edge-frequency distribution of the CW spectra as described in Section **4.5.1**, is shown in Figure **4.16**. The projected silhouette in the plane-of-sky appears non-circular, exhibiting a roughly triangular shape with three smoothly rounded lobes and a mild concavity along the lower-left boundary. The center of mass, shown as a red cross, is offset from the geometric centroid of the silhouette toward the upper-right quadrant. Under the assumption of a homogeneous internal density distribution, this offset is consistent with the asymmetric mass distribution implied by the non-circular projected shape. In principle, however, a similar offset could also arise from a non-uniform internal density distribution even in a more symmetric body; radar observations alone cannot distinguish between these two cases.

The semi-axes of the convex hull, corrected for the subradar latitude $\delta = 16.5°$, are $a = 185$ m and $b = 191$ m, yielding an axis ratio $b/a = 1.03$. The near-equality of a and b is physically plausible: a body with $a/b \approx 1$ appears nearly symmetric in projection regardless of rotational phase. The systematic departure from circular symmetry, most evident in the flattened lower boundary and the protruding upper-right lobe, suggests genuine non-sphericity at the decametric scale.

It should be recalled that the convex hull is a projection of the equatorial cross-section: it represents the maximum extent of the body as seen from the subradar direction, and provides no information about the polar dimension or the true three-dimensional shape. The convex hull is therefore an upper bound on the equatorial dimensions and any concavities on the actual surface are projected inside the hull boundary. The recovered silhouette is best interpreted as constraining the equatorial aspect ratio and the approximate azimuthal distribution of the bulk material.

4.5.3 3D rotational ellipsoid

The natural first extension of the convex hull model described in Section **4.5.1** is to fit a triaxial ellipsoid to the full set of observed CW spectra, which means to exploit the information contained in the shape of the Doppler power spectra, and not just considering their edge frequencies. The ellipsoid model is the simplest closed, three-dimensional figure consistent with a non-zero polar extent, and its four parameters (three semi-axes and the subradar latitude) are directly interpretable in terms of physical dimensions and observing geometry.

In the pipeline described below, the a and b semi-axes are taken as fixed values from the convex hull reconstruction, while the subradar latitude has been inherited from external observations, leaving a single free parameter to be optimized against the data. This progressive constraint strategy is a deliberate design choice: it reduces the dimensionality of the optimization problem, limits the degeneracies, and ensures that each level of the hierarchy builds on and is consistent with the results of the previous one.

The geometry of the problem is shown in Figure **4.17**. The asteroid is modeled as a triaxial ellipsoid with semi-axes a,b,c aligned with the axes of a body-fixed rotating coordinate system (X',Y',Z'), where Z' is the spin vector. The observer is placed at a large distance along the negative X-direction of the fixed inertial frame, so the radar line-of-sight is $\hat{v} = (-1, 0, 0)^T$.

The orientation of the spin vector relative to the line-of-sight is described by the subradar latitude δ (in degrees), here expressed in degrees, defined as the complement of the aspect angle $\alpha = 90° - \delta$ between the spin vector and the line-of-sight. The spin axis in the fixed frame is:

$$\hat{Z'} = (\cos\alpha, 0, \sin\alpha)^T = (\sin\delta, 0, \cos\delta)^T \tag{4.26}$$

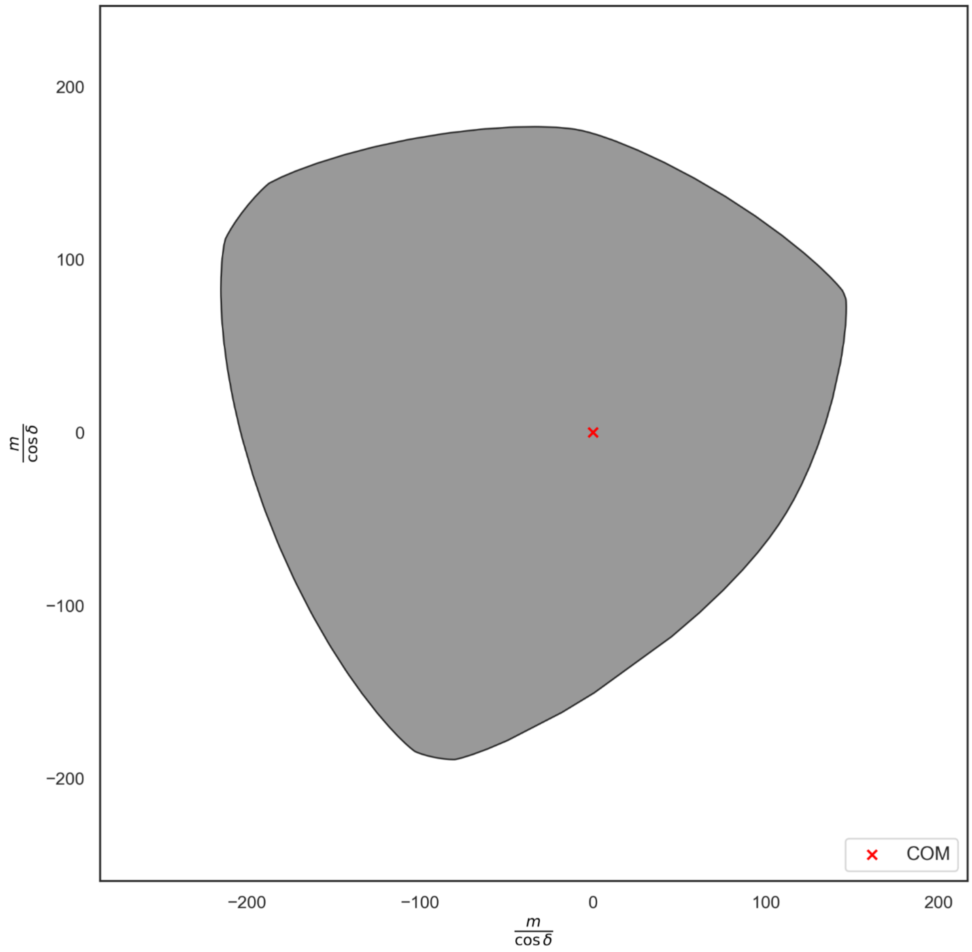

Figure 4.16: Pole-on projection of the 2005 LW3 convex envelope. The red cross marks the rotation center.

A right-handed orthonormal frame $(X_{\text{rot}}, Y_{\text{rot}}, Z_{\text{rot}})$ is defined with $Z_{\text{rot}} = \hat{Z}'$. X_{rot} lies in the plane spanned by $\hat{v}$ and $\hat{Z}'$ and orthogonal to $\hat{Z}'$, and $Y_{\text{rot}} = Z_{\text{rot}} \times X_{\text{rot}}$. The transformation matrix $\mathbf{R}_0 = [X_{\text{rot}}|Y_{\text{rot}}|Z_{\text{rot}}]$ maps from the rotating frame to the fixed frame at the reference epoch. At time t corresponding to rotational phase angle ϕ, the body rotates about $\hat{Z}'$ with angular velocity $\omega = 360°/P$, so the total transformation from body-fixed to fixed frame is:

$$\mathbf{R}_{\text{total}} = \mathbf{R}_0 \cdot \mathbf{R}_z(\omega t) \tag{4.27}$$

where $\mathbf{R}_z(\omega t)$ is the rotation matrix about Z' by angle $\beta = \omega t$.

The ellipsoid surface is discretized using an icosphere method: starting from a regular icosahedron (20 triangular faces, 12 vertices) and iteratively subdividing each triangle into four sub-triangles whose vertices are projected back onto the unit sphere, a set of $N_{\text{pts}} = 10000$ points approximately uniformly distributed on the unit sphere is generated. Each point (x_s, y_s, z_s) on the unit sphere is then mapped to the ellipsoid surface by coordinate scaling: $(X_l, Y_l, Z_l) = (ax_s, by_s, cz_s)$. This method provides a much more uniform coverage of the surface than traditional spherical grid sampling (based on longitude and latitude), which over-samples the polar regions and under-samples the equatorial band.

The outward unit normal at each surface point, in the body-fixed frame, is:

$$\hat{N}_l = \frac{1}{\|\cdot\|}\left(\frac{X_l}{a^2}, \frac{Y_l}{b^2}, \frac{Z_l}{c^2}\right) \tag{4.28}$$

Both the position vectors and the normals are rotated to the fixed frame via $\mathbf{R}_{\text{total}}$. The radial velocity of each surface point with respect to the observer is the component of the surface velocity along the line-of-sight (LOS):

$$v_r = -(\omega \times \mathbf{r}_{\text{fixed}}) \cdot \hat{v} \tag{4.29}$$

where $\omega = \mathbf{R}_{\text{total}}(0, 0, \omega)^T$ is the angular velocity vector in the fixed frame and $\mathbf{r}_{\text{fixed}}$ is the position of the surface point. Only points on the hemisphere facing the observer (those with $\hat{N}_{\text{fixed}} \cdot \hat{v} > 0$) contribute to the echo. Each visible point contributes a Doppler shift as in Equation **2.76**.

The scattered power from each visible surface element is modeled using a Lambertian scattering law, in which the reflected intensity is proportional to $\cos^2\theta_i$, where θ_i is the angle between the surface normal and the radar LOS. This formula is the standard first-order approximation used in planetary radar echo modeling [212] when no detailed surface scattering characterization is available.

The simulated Doppler power spectrum is generated by binning the scattered power of each visible surface point into the frequency grid defined by the observed spectrum, using a triangular distribution kernel over three adjacent bins to reduce aliasing artifacts from the discrete frequency grid. The resulting spectrum $S^{(i)}_{\text{sim}}(f_j, c)$ is a physically motivated forward model that differs fundamentally from the empirical Jurgens model of Eq. **4.7** used in the spectral analysis pipeline. The Jurgens model is a parametric fit to the spectral profile shape with no geometric interpretation, designed to estimate the echo bandwidth. In contrast, S_{sim} is derived directly from the three-dimensional geometry of the rotating ellipsoid: the power distribution across Doppler frequencies depends explicitly on the polar semi-axis c, which controls the out-of-plane extent of the body and hence the fraction of the surface visible at each rotational phase. This makes S_{sim} sensitive to the full three-dimensional shape of the asteroid, not just its projected equatorial bandwidth.

The semi-axes a and b are corrected for the subradar latitude factor $\cos\delta$:

$$a_{\text{phys}} = \frac{a_{\text{hull}}}{\cos\delta}, \quad b_{\text{phys}} = \frac{b_{\text{hull}}}{\cos\delta} \tag{4.30}$$

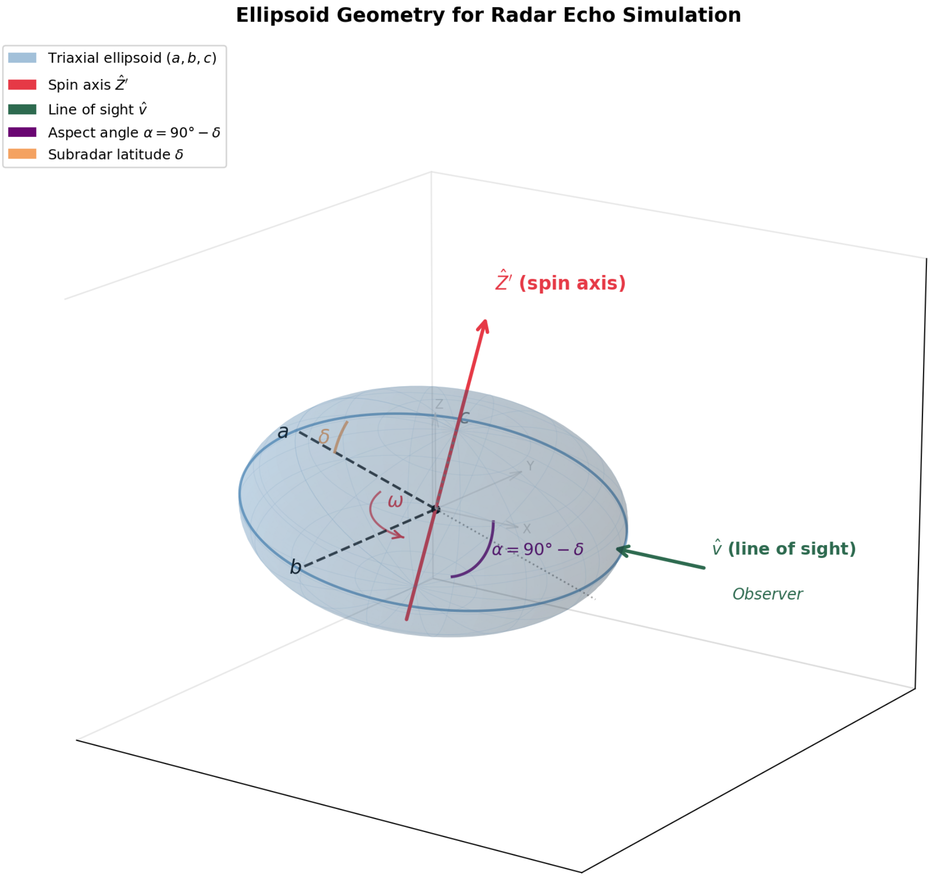

Figure 4.17: Sketch of the reference ellipsoid geometry

The subradar latitude δ is constrained by delay-Doppler radar observations of 2005 LW3 performed during the same close approach [206]. The value $\delta = 16.5°$ derived from these observations is adopted as a fixed external constraint throughout the ellipsoid fitting. We stressed the fact that δ is the only parameter in this reconstruction that is not determined from the CW radar data themselves. If it were unknown, we could have used a_{hull} and b_{hull} for the analysis as the maximum dimensions of the object.

The single free parameter of the optimization is therefore the polar semi-axis c, which determines the extent of the ellipsoid along the spin vector and is not constrained by the two-dimensional convex hull. The optimization minimizes a chi-squared objective function:

$$\chi^2(c) = \sum_{i=1}^{L} \sum_{j \in B_i} \frac{\left[S_{obs}^{(i)}(f_j) - S_{sim}^{(i)}(f_j, c)\right]^2}{\sigma_{i,j}^2} \tag{4.31}$$

where the outer sum runs over all L spectra in the dataset, and the inner sum runs over frequency channels j, within the echo bandwidth B_i of each i-th spectrum. The minimization is performed using the Levenberg–Marquardt algorithm [213] via a MATLAB routine, with $c \in [100, 1000]$ m and initialized at $c_0 = 120$ m. The computation is parallelized across spectra using MATLAB's Parallel Computing Toolbox.

The quality of the fit is assessed through the reduced chi-squared:

$$\chi_\nu^2 = \frac{\chi^2}{N_{\text{data}} - 1} \tag{4.32}$$

where N_{data} is the total number of spectral channels across all spectra. Values of χ_ν^2 close to unity indicate that the ellipsoid model describes the data at the level of the noise; values significantly greater than unity indicate systematic residuals that motivate the more flexible shape models of the subsequent reconstruction levels. Conversely, values of χ_ν^2 significantly less than unity may indicate either overfitting in the optimization (where the model has been tuned to follow noise rather than signal) or an overestimation of the uncertainties assigned to the input data.

The polar semi-axis c estimated by this procedure, combined with the equatorial semi-axes a and b from the convex hull, yields a complete triaxial ellipsoid describing the large-scale three-dimensional figure of 2005 LW3, with the axis ratios a/b and a/c characterize the equatorial elongation and the polar flattening, respectively. The ellipsoid model has two fundamental limitations. First, it imposes strict convexity and smoothness: any deviation from an ellipsoidal shape, e.g. concavities, elongated lobes, large-scale topographic features, is invisible to this model. Second, the Lambertian scattering law is a simplified description of the true surface scattering, and the actual radar echo spectrum will differ from the simulated one even for a perfectly ellipsoidal body if the surface has heterogeneous roughness or scattering properties. These limitations motivate the more general shape reconstruction approaches described in Sections **4.5.5** and beyond, which progressively relax the shape constraints while retaining the physical geometry established at this level.

4.5.4 Ellipsoid model results

The best-fit triaxial ellipsoid has semi-axes $a = 185$ m, $b = 191$ m, and $c = 226$ m. The polar semi-axis c exceeds the equatorial semi-axes by approximately 20%, yielding axis ratios $c/a = 1.22$ and $c/b = 1.18$. The body is elongated along its spin axis, a morphology sometimes described as oblate-to-prolate transitional, though a definitive classification requires independent constraints on the spin pole orientation that are not available from CW data alone. Since the subradar latitude is low, the observer was close to the equatorial plane during the observations, which maximizes the sensitivity of the CW spectra to the equatorial dimensions while reducing the leverage on the polar extension c. The recovered value of c should therefore be regarded as carrying larger systematic uncertainty than a and b.

The spectral fits for six representative rotational phases uniformly sampled across the full 360° are shown in Figure **4.18**. The simulated spectra (red dashed) reproduce the overall shape and bandwidth of the observed echoes (black solid) well at all phases, capturing both the Doppler edge positions and the broad curvature of the spectral envelopes. The reduced chi-squared is $\chi^2_\nu = 1.347$. A value of χ^2_ν moderately above unity is expected for a smooth ellipsoid model applied to echoes from a body whose surface shows genuine small-scale topography: the model accounts for the overall shape but cannot reproduce the rapid spectral fluctuations visible in the observed echoes, which reflect coherent backscattering from individual surface facets rather than a smooth reflector. The systematic residuals are most prominent near the spectral peaks, where the observed spectra show sharp spikes that a smooth ellipsoid cannot generate by definition.

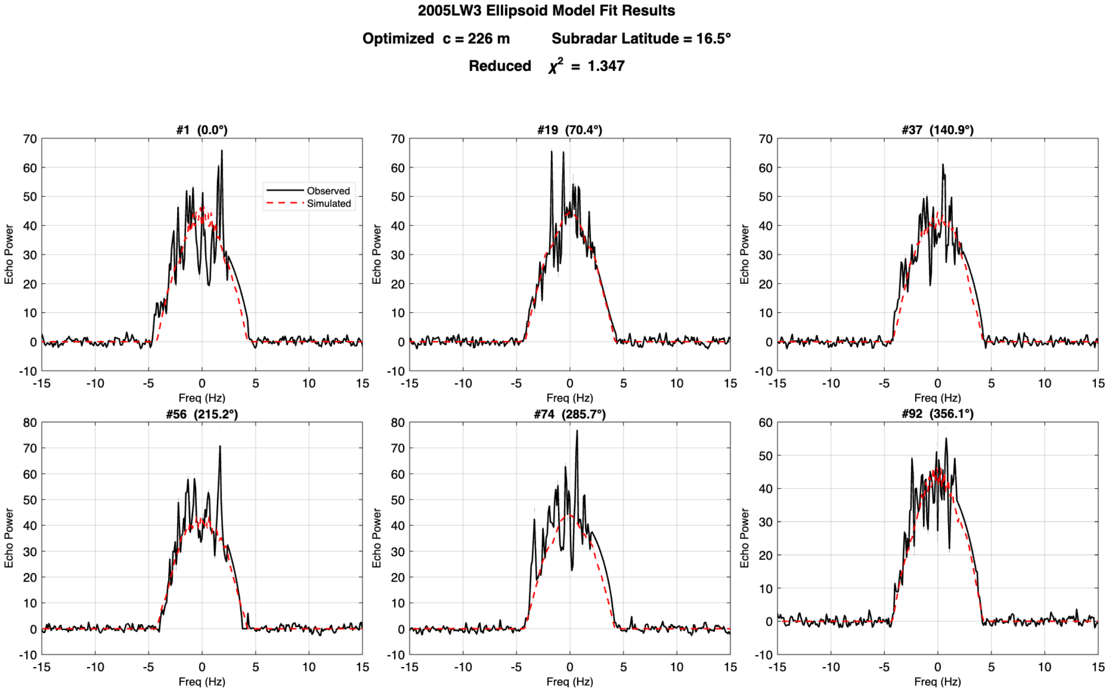

Figure 4.18: Simulated ellipsoid spectra (dotted red line) vs real data (black line). Frequency resolution: 0.1 Hz, integration time: 10 minutes.

The three-dimensional ellipsoid is shown in Figure **4.19**, color-coded by surface normal angle relative to the observer direction. The sky-plane projection confirms that the equatorial cross-section is nearly circular, which is consistent with the b/a ratio derived before. The visible hemisphere at rotational phase $\phi = 45°$ is dominated by a broad central region with small surface normal angles, since the specular region contributes most strongly to the echo by construction.

4.5.5 Radial perturbation model

The ellipsoid model of Section **4.5.3** constrains the large-scale three-dimensional figure of 2005 LW3 but is by construction incapable of describing any deviation from a smooth, ellipsoidal surface. Real asteroids, particularly small bodies in the size range of a few hundred meters, are expected to have

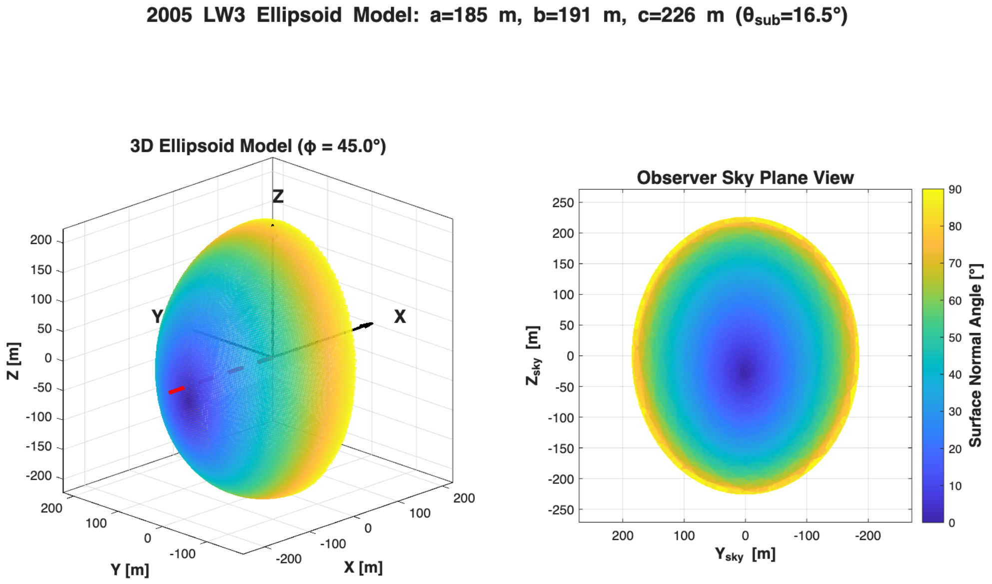

Figure 4.19: Sketch of the reference ellipsoid

significant surface topography, such as craters, ridges, flat facets, and concavities, whose signatures may be detectable in the shape of the Doppler power spectrum even at the signal-to-noise levels available in this dataset, at least the most prominent ones, if present. The radial perturbation model addresses this limitation by introducing a set of localized, independent deviations from the reference ellipsoid surface, optimized to minimize the residuals between observed and simulated spectra.

The surface of the perturbed body is defined as a radial displacement field superimposed on the reference ellipsoid. As shown in Figure **4.20**, the body surface is divided into $N_\theta \times N_\phi = 8 \times 16 = 128$ rectangular patches in spherical coordinates (θ, ϕ), where $\theta \in [0°, 180°]$ is the co-latitude measured from the north pole and $\phi \in [0°, 360°]$ is the longitude. The patch centers are placed at uniformly spaced values:

$$\theta_i = \theta_{\min} + \frac{i - \frac{1}{2}}{N_\theta}(\theta_{\max} - \theta_{\min}), \quad i = 1, ..., N_\theta \tag{4.33}$$

$$\phi_j = \frac{(j-1) \cdot 360°}{N_\phi}, \quad j = 1, ..., N_\phi \tag{4.34}$$

A small margin $\theta_{\text{margin}} = 0.1$ rad $\approx 5.73°$ is taken at both poles to avoid the numerical singularities that arise at $\theta = 0°$ and $\theta = 180°$, where the surface normals become degenerate and the Doppler velocity of surface points vanishes exactly. Each patch carries an independent scalar perturbation parameter $\delta r_{ij} \in [-25, +25]$ m, which displaces the surface radially by that amount relative to the reference ellipsoid at all points failing within the patch. The full parameter vector $\delta\mathbf{r} \in \mathrm{R}^{128}$ collects all patch perturbations. Given a surface point at co-latitude θ and longitude ϕ, its patch is identified by nearest-bin assignment: the patch index (i, j) is found by normalizing the angular coordinates into $[0, 1]$ and computing $i = [N_\theta \cdot \theta_{\text{norm}}], j = [N_\phi \cdot \phi_{\text{norm}}]$, with values constrained within the admissible interval. Each icosphere point from the surface discretization is assigned to the patch that contains it, and the position of that point is then displaced radially by the patch perturbation:

$$\mathbf{r}_{\text{pert}} = \mathbf{r}_{\text{base}} + \delta r_{ij} \cdot \hat{r}_{\text{base}} \tag{4.35}$$

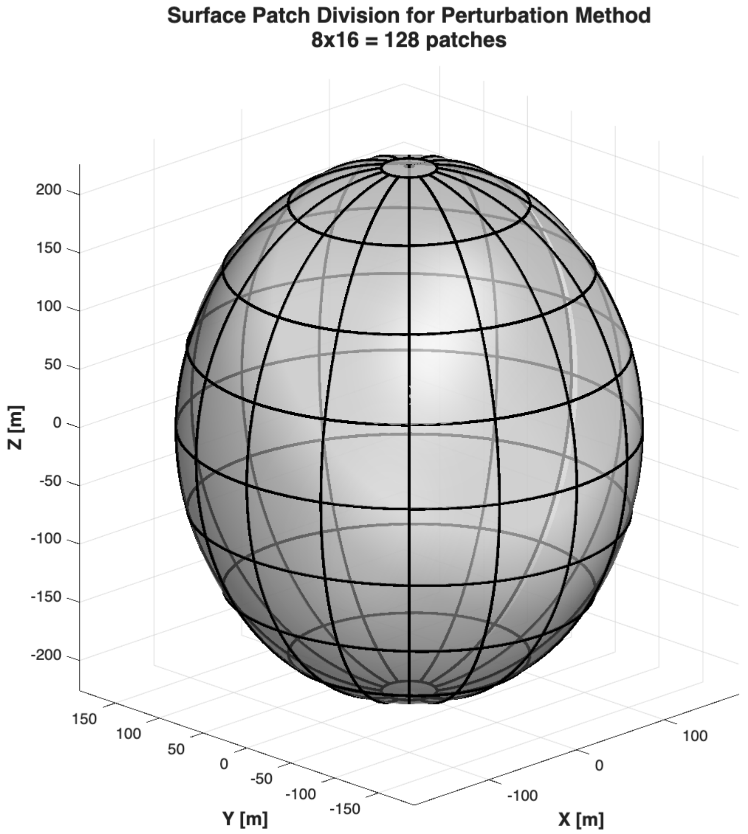

Figure 4.20: Radial surface patch division of the reference ellipsoid

where $\hat{r}_{\text{base}} = \hat{\mathbf{r}}_{\text{base}}/|\hat{\mathbf{r}}_{\text{base}}|$ is the unit radial vector at that point on the base ellipsoid. The perturbation is therefore purely radial and does not alter the angular position of any surface point. The introduction of radial perturbations poses a difference in the calculation of the surface normals of the smooth ellipsoid model. In that case, the outward normal at a point is described in Equation **4.28**. For the perturbed surface, this formula is no longer valid because the surface has sharp discontinuities at patch boundaries, and the normal direction changes discontinuously between adjacent patches. Thus, the normals are computed numerically using the gradient of the implicit surface function $F(\mathbf{r}) = 0$ that defines the perturbed shape. The perturbed surface is described implicitly by:

$$F(x,y,z) = \left(\frac{x}{a}\right)^2 + \left(\frac{y}{b}\right)^2 + \left(\frac{z}{c}\right)^2 - \left(\frac{r_{\text{target}}(\theta,\phi)}{r_{\text{ellipsoid}}(\theta,\phi)}\right)^2 = 0 \tag{4.36}$$

where $r_{\text{ellipsoid}}(\theta,\phi)$ is the radial distance from the origin to the unperturbed surface in direction (θ,ϕ):

$$r_{\text{ellipsoid}}(\theta,\phi) = \left(\frac{\cos^2\phi\sin^2\theta}{a^2} + \frac{\sin^2\phi\sin^2\theta}{b^2} + \frac{\cos^2\theta}{c^2}\right)^{-\frac{1}{2}} \tag{4.37}$$

and $r_{\text{target}}(\theta,\phi) + \delta r_{ij}$ is the target radial distance after perturbation. The gradient ∇F is evaluated numerically at each surface point central finite differences [214]. If the gradient magnitude vanishes, which can occur at patch boundaries where the perturbation field is discontinuous, the normal falls back to the ellipsoidal formula for that point.

The Doppler spectrum generation for the perturbed body follows the same procedure as for the smooth ellipsoid. The only differences are that the surface point positions and normals now depend on the perturbation field $\delta\mathbf{r}$ rather than being fixed by the ellipsoidal geometry. Additionally, the optimization minimizes the same weighted residual objective function. In this case, the three ellipsoid semi-axes (a,b,c) and the subradar latitude δ are held fixed at the values determined in the previous reconstruction steps. The 128 perturbation parameters δr_{ij} are initialized to zero, which corresponds to the smooth reference ellipsoid, and bounded to $|\delta r_{ij}| \leq 25$ m throughout the optimization.

The radial perturbation model provides a direct, model-independent estimate of the surface topography of 2005 LW3 at the angular resolution of the patch grid, that is approximately 20° in latitude and 22.5° in longitude, corresponding to length scales of order $60 - 80$ m on the surface. Features smaller than one patch (below this resolution) are not resolved by the model. This resolution was intentionally selected to capture only large-scale concavities or surface perturbations of the asteroid. Detecting smaller-scale features with a Doppler-only analysis is inherently challenging, as noise and reconstruction artifacts may be misinterpreted as genuine surface structures. Positive values of δr_{ij} indicate regions where the surface extends beyond the reference ellipsoid, such as ridges, bulges, or elongated lobes. In contrast, negative values indicate depression or flat facets.

Two important limitations arise from this implementation. First, the patch assignment is a nearest-bin step function, so the perturbation field is piecewise constant and introduces sharp discontinuities at patch boundaries. This means that the perturbed surface normals computed numerically at those boundaries are artifacts of the discretization rather than physical features of the shape.

Second, the patch parameterization treats each patch as an independent degree of freedom with no spatial regularization: nothing in the optimization prevents adjacent patches from having large, opposite perturbations that cancel in the integrated spectrum but produce unphysical surface morphologies. In principle, a regularization penalty on the spatial gradient of $\delta \mathbf{r}$ could be added to suppress such solutions, but this was not included in the current implementation to leave the model maximally unconstrained by prior assumptions on the surface smoothness.

4.5.6 Radial perturbation model results

The radial perturbation model described in Section **4.5.5**, with 128 free patch amplitudes on an 8×16 angular grid, yields a best-fit perturbation field with standard deviation $\sigma_{\delta r} = 4.8$ m and full range $[-14.1, +14.1]$ m. The maximum departure from the reference ellipsoid is thus ~ 14 m, corresponding to 7.5% of the minor equatorial semi-axis a. The spectral fits are shown in Figure **4.21** and the recovered three-dimensional shape is presented in Figure **4.22**.

The fit quality improves with respect to the ellipsoid model: χ^2_ν decreases from 1.347 to 1.186. The remaining excess above unity reflects residual structure in the observed spectra that the 8×16 patch resolution cannot capture, either because it corresponds to sub-patch-scale topography or because it arises from coherent speckle noise in the echo. Examining the three-dimensional shape (Figure **4.22**) the perturbed ellipsoid shows a concentration of positive perturbations near the north polar region, visible in both the standard view and the side view (YZ plane). This feature is the most prominent departure from the reference ellipsoid and corresponds to rotational phases where the observed spectra show the broadest deviation from the smooth ellipsoid model. A secondary region of positive perturbation is visible in the front view (XZ plane) at mid-latitudes, while the equatorial regions show predominantly negative perturbations distributed roughly symmetrically in longitude. The top view (XY plane) confirms that the perturbation pattern preserves the near-circularity of the equatorial cross-section established by the convex hull.

These morphological features must be interpreted with caution, as explained in Section **4.5.5**. The

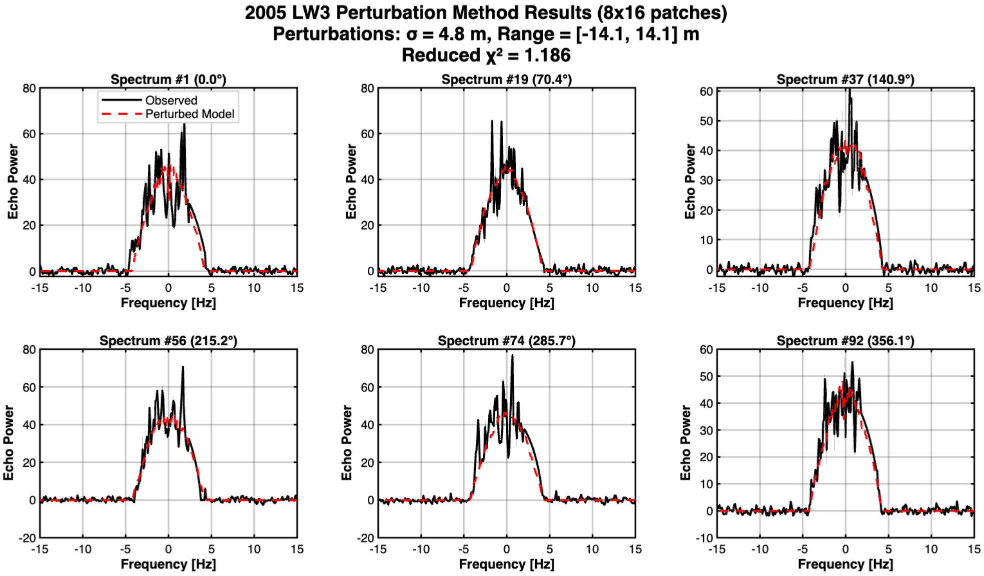

Figure 4.21: Simulated radial perturbed ellipsoid spectra (dotted red line) vs real data (black line). Frequency resolution: 0.1 Hz, integration time: 10 minutes.

CW data constrain the equatorial projection most directly, and perturbations at high latitudes, such as the north polar features, are substantially less constrained than equatorial ones. The high-latitude topography may reflect the model absorbing spectral noise rather than real surface features (see Section **4.5.10**). Nevertheless, the consistent improvement in χ^2_ν across all phases indicates that the equatorial perturbation carries real information, and the amplitude scale $\sigma_{\delta r} = 4.8$ m is a physically meaningful estimate of the RMS surface relief at the spatial scale of the patch grid.

4.5.7 Vectorial perturbation model

The radial perturbation model of Section **4.5.5** displaces each surface patch along the local outward normal, effectively adding or removing material in the radial direction. A natural generalization is to allow the surface displacement to have an arbitrary orientation, described by three independent components in the local spherical coordinate frame. This vectorial perturbation model introduces, for each surface patch, a displacement vector:

$$\delta(\theta, \phi) = \delta r \cdot \hat{r} + \delta\theta \cdot \hat{\theta} + \delta\phi \cdot \hat{\phi} \tag{4.38}$$

where $\hat{r}, \hat{\theta}, \hat{\phi}$ are the unit vectors of the local spherical coordinate frame: $\hat{r}$ is the outward radial direction, $\hat{\theta}$ is the co-latitudinal direction (pointing toward the equator from the north pole), and $\hat{\phi}$ is the azimuthal direction (pointing toward increasing longitude). The three scalar perturbation amplitudes $(\delta r, \delta\theta, \delta\phi)$ are independently optimized for each of the $N_\theta \times N_\phi = 128$ surface patches, yielding a total of 384 free parameters.

The surface grid, patch assignment, and Doppler spectrum generation are identical to what has been described above. The only difference is in how the position of each surface point is computed: instead of a purely radial displacement, the full three-component displacement vector is applied to each icosphere point, and the surface normals are recomputed numerically using central finite differences of the implicit surface function.

The local spherical unit vectors are computed analytically from the Cartesian coordinates of each

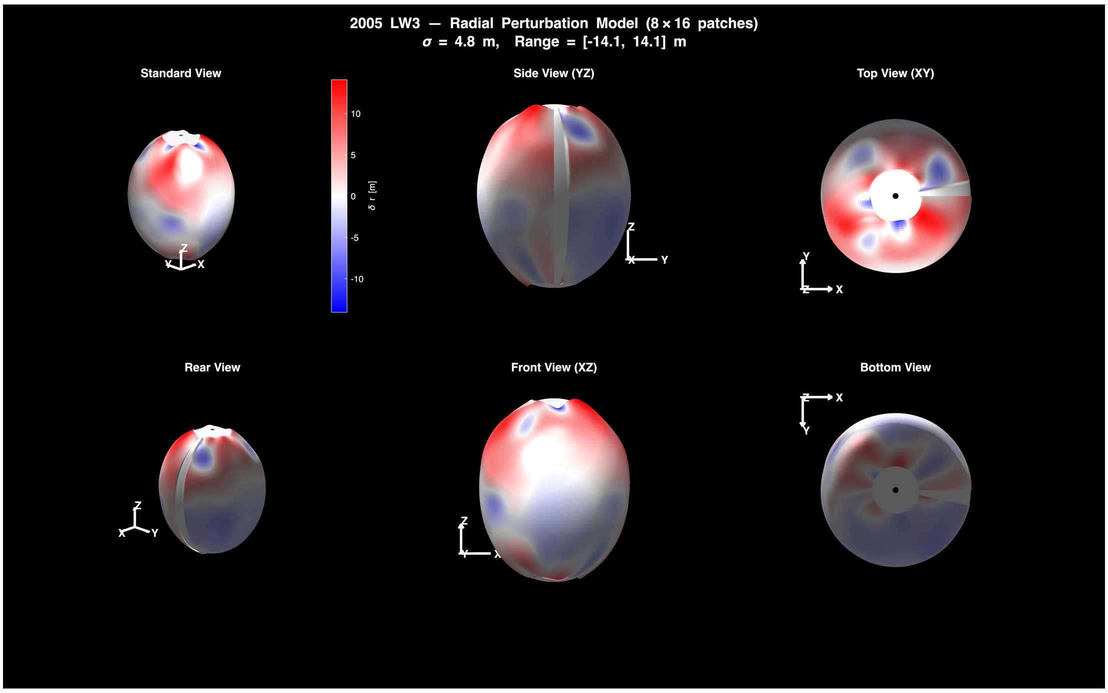

Figure 4.22: Three-dimensional shape of the radial perturbed reference ellipsoid. The surface is color-coded according to the radial perturbation δr at each vertex relative to the reference ellipsoid: blue indicates regions where the reconstructed surface lies below the reference ellipsoid, corresponding to depressions, concavities, or flat facets; white corresponds to near-zero perturbation, i.e. surface elements that locally coincide with the smooth ellipsoid, and red indicates regions where the surface extends beyond the reference ellipsoid, corresponding to ridges, bulges, or elongated lobes.

surface point. For a point *(x,y,z)*, with radial distance $r = \sqrt{x^2 + y^2 + z^2}$ and equatorial projection $\rho = \sqrt{x^2 + y^2}$:

$$\begin{aligned}
\hat{r} &= \left(\frac{x}{r}, \frac{y}{r}, \frac{z}{r}\right)^T = (\sin\theta\cos\phi, \sin\theta\sin\phi, \cos\theta)^T \\
\hat{\theta} &= \left(\frac{xz}{r\rho}, \frac{yz}{r\rho}, -\frac{\rho}{r}\right)^T = (\cos\theta\cos\phi, \cos\theta\sin\phi, -\sin\theta)^T \qquad (4.39) \\
\hat{\phi} &= \left(-\frac{y}{\rho}, \frac{x}{\rho}, 0\right)^T = (-\sin\phi, \cos\phi, 0)^T
\end{aligned}$$

The bounds on the perturbation components are $|\delta r| \leq 25$ m , $|\delta\theta| \leq 0.1$ rad $\approx 5.73°$ and $|\delta\phi| \leq 0.1$ rad $\approx 5.73°$, with all parameters initialized to zero. The optimization algorithm is the same as before, with χ^2_ν evaluated with $N_{\text{pert}} = 384$ degrees of freedom subtracted.

Before discussing the results in the next section, it is important to assess the physical content of the two tangential components $\delta\theta$ and $\delta\phi$ considered in the context of CW radar observations.

The co-latitudinal component $\delta\theta$ displaces surface material toward higher or lower latitudes, modifying the distribution of the asteroid's surface along the spin axis. However, in a CW radar observation, the observed power spectrum is determined entirely by the distribution of surface area as a function of Doppler frequency, which in turn depends on the radial velocity component of each surface element. Consequently, purely latitudinal displacements ($\delta\theta \neq 0, \delta r = \delta\phi = 0$) modify the out-of-plane shape of the body without changing its equatorial projection, and therefore have a negligible effect on the CW echo spectrum at first order. Hence, the $\delta\theta$ parameters are nearly degenerate with respect to CW data.

The azimuthal component $\delta\phi$ displaces surface material along the direction of rotation, redistributing it among different Doppler frequency bins at any given rotational phase. This does have a direct effect on the spectrum shape, equivalent to locally shifting the angular position of surface features. However, $\delta\phi$ perturbations introduce a near-degeneracy with the rotational phase assignment: a coherent azimuthal shift of all patches is nearly indistinguishable from a small error in the rotation period. Only spatially incoherent $\delta\phi$ patterns produce new real spectral structure not described by a rotational phase shift.

In summary, of the 384 free parameters of this method, 128 carry the same physical content as the previous implementation (δr), 128 provide a modest additional flexibility in redistributing spectral power azimuthally ($\delta\phi$), and the last 128 are largely unconstrained by CW data. These consideration have been taken into the following, when we discuss the products of this reconstruction algorithm.

4.5.8 Vectorial perturbation model results

The vectorial pertubation model, with 384 free parameters encoding radial (δr), co-latitudinal ($\delta\theta$), and azimuthal ($\delta\phi$) displacements on the same 8×16 patch grid, produces a best-fit solution in which the perturbation field is numerically negligible: the radial component standard deviation is $\sigma_{\delta r} \approx 0$ m with a full range of $[-0.1, +0.2]$ m, and the recovered three-dimensional shape is virtually indistinguishable from the smooth reference ellipsoid. The spectral fits are shown in Figure **4.23** and the recovered shape in Figure **4.24**.

The fit metric confirm the failure of the model: $\chi^2_\nu = 1.428$, worse than the smooth ellipsoid ($\chi^2_\nu = 1.347$) and substantially worse than the radial perturbation model ($\chi^2_\nu = 1.186$). The model with three times as many parameters as the radial perturbation model performs worse than the model with no free parameters beyond the reference ellipsoid.

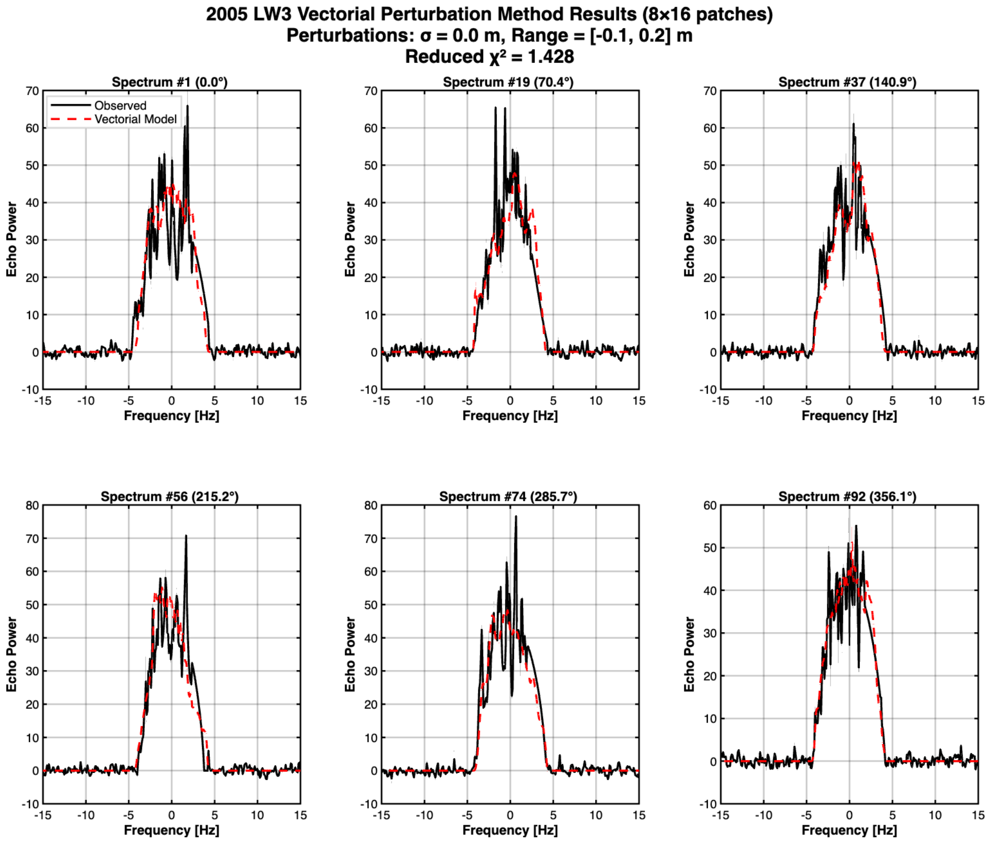

Figure 4.23: Simulated vectorial perturbed ellipsoid spectra (dotted red line) vs real data (black line). Frequency resolution: 0.1 Hz, integration time: 10 minutes.

This outcome is not a numerical failure of the optimizer but an empirical confirmation of the degeneracy analysis presented in Section **4.5.7**. Of the 384 parameters, the 128 radial components δr_i carry the same physical information as the previous method and are constrained by the same CW observables. Conversely, the 128 co-latitudinal components $\delta\theta_i$ are essentially unconstrained: a displacement along $\hat{\theta}$ moves a surface point toward or away from the equatorial plane, which does not alter the projected area at any Doppler frequency bin and therefore produces no spectral signature in the CW data. The 128 azimuthal components $\delta\phi_i$ are nearly degenerate with the rotational phase and contribute only weakly and indirectly to the spectral shape. The optimizer therefore operates in a 384-dimensional parameter space of which approximately 256 dimensions are flat (the cost function has no gradient along these directions), leaving the solver unable to find a descent direction that meaningfully improves the fit. The consequence is that the solver stagnates near the initial condition (the smooth ellipsoid), with the 256 unconstrained parameters remaining at zero and the 128 radial parameters unable to converge because the gradient landscape along those directions is corrupted by the flat tangential subspace. The resulting $\chi^2_\nu > 1.347$ reflects the additional degrees of freedom subtracted from the denominator without any corresponding improvement in the numerator.

In conclusion, the information content of the CW dataset is insufficient to constrain a full vectorial perturbation model on this patch grid, and adding tangential degrees of freedom actively harms the optimization rather than improving it. This is an instance of a well-known phenomenon in ill-posed inverse problems: the introduction of unconstrained parameters degrades the recovery of the constrained ones by creating flat directions that destabilize the solver. The present result quantifies its effect on a concrete radar dataset.

This finding strengthens the interpretation of the radial perturbation method results: the perturbation amplitudes recovered by the radial model are not an artifact of an underfitted parameterization but represent the genuine limit of what CW radar can resolve about the shape of 2005 LW3.

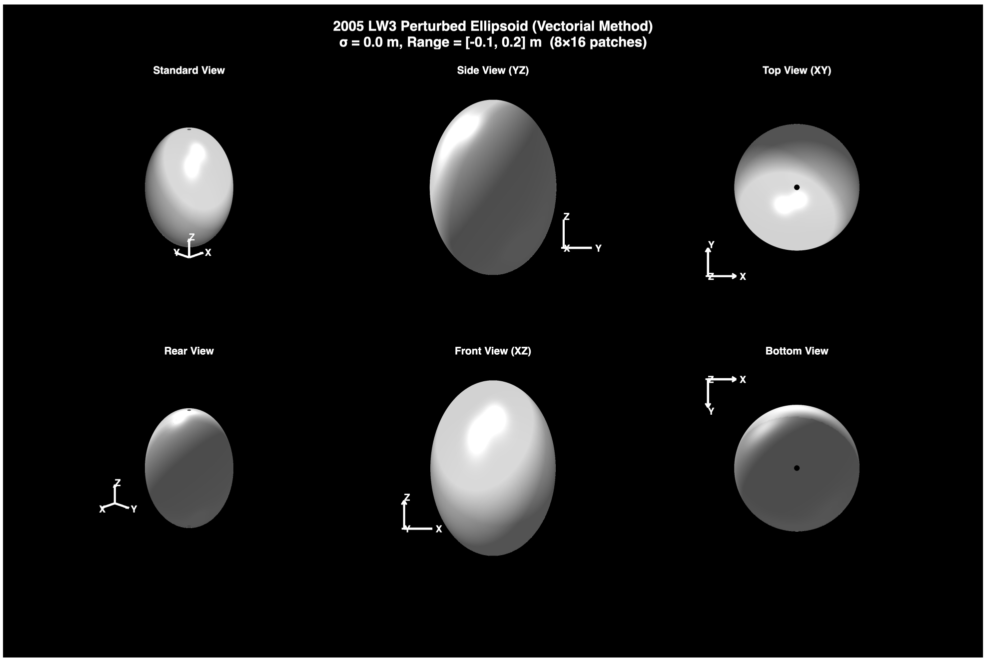

Figure 4.24: Three-dimensional shape of the vectorial perturbed reference ellipsoid.

4.5.9 Adaptive vertex model with local perturbations

The perturbation models described in the previous sections divided the surface into a fixed, uniform grid of patches whose resolution is identical everywhere on the body, regardless of where the data contain meaningful shape information. For a CW dataset, the spectral constraints are not spatially uniform, since certain rotational phases and equatorial longitudes produce stronger, more distinctive echo features than others, and the associated shape information is localized to specific surface regions. A uniform patch grid, therefore, can not adapt to this anisotropic behavior of Doppler spectra.

The adaptive vertex model addresses this issue by representing the surface as triangulated mesh $\mathcal{M} = \{V, F\}$, where $V = \{\mathbf{v}_i\}_{i=1}^{N}$ is the set of vertex positions in $\mathbb{R}^3$ and F is the set of triangular faces. Each vertex carries a scalar radial perturbation δr_i, defined as its displacement from the reference ellipsoid along the local outward normal $\hat{n}_i$:

$$\mathbf{v}_i = \mathbf{v}_i^0 + \delta r_i \, \hat{n}_i \tag{4.40}$$

where $\mathbf{v}_i^0$ is the position of vertex i on the unperturbed ellipsoid with semi-axes (a,b,c) determined from the ellipsoid model reconstruction. The vertex count N is not fixed: new vertices are introduced iteratively in regions where the data-model residuals are largest, so that the effective spatial resolution of the model tracks the information content of the dataset. The algorithm is summarized in the flowchart of Figure **4.25**. The initial mesh is constructed according to the icosphere method of Section **4.5.3**: starting with 12 vertices and 20 triangular faces, it is recursively subdivided. Each triangular face is split into four smaller triangles by introducing the midpoints of its three edges, which are then projected onto the unit sphere until the vertex count exceeds a

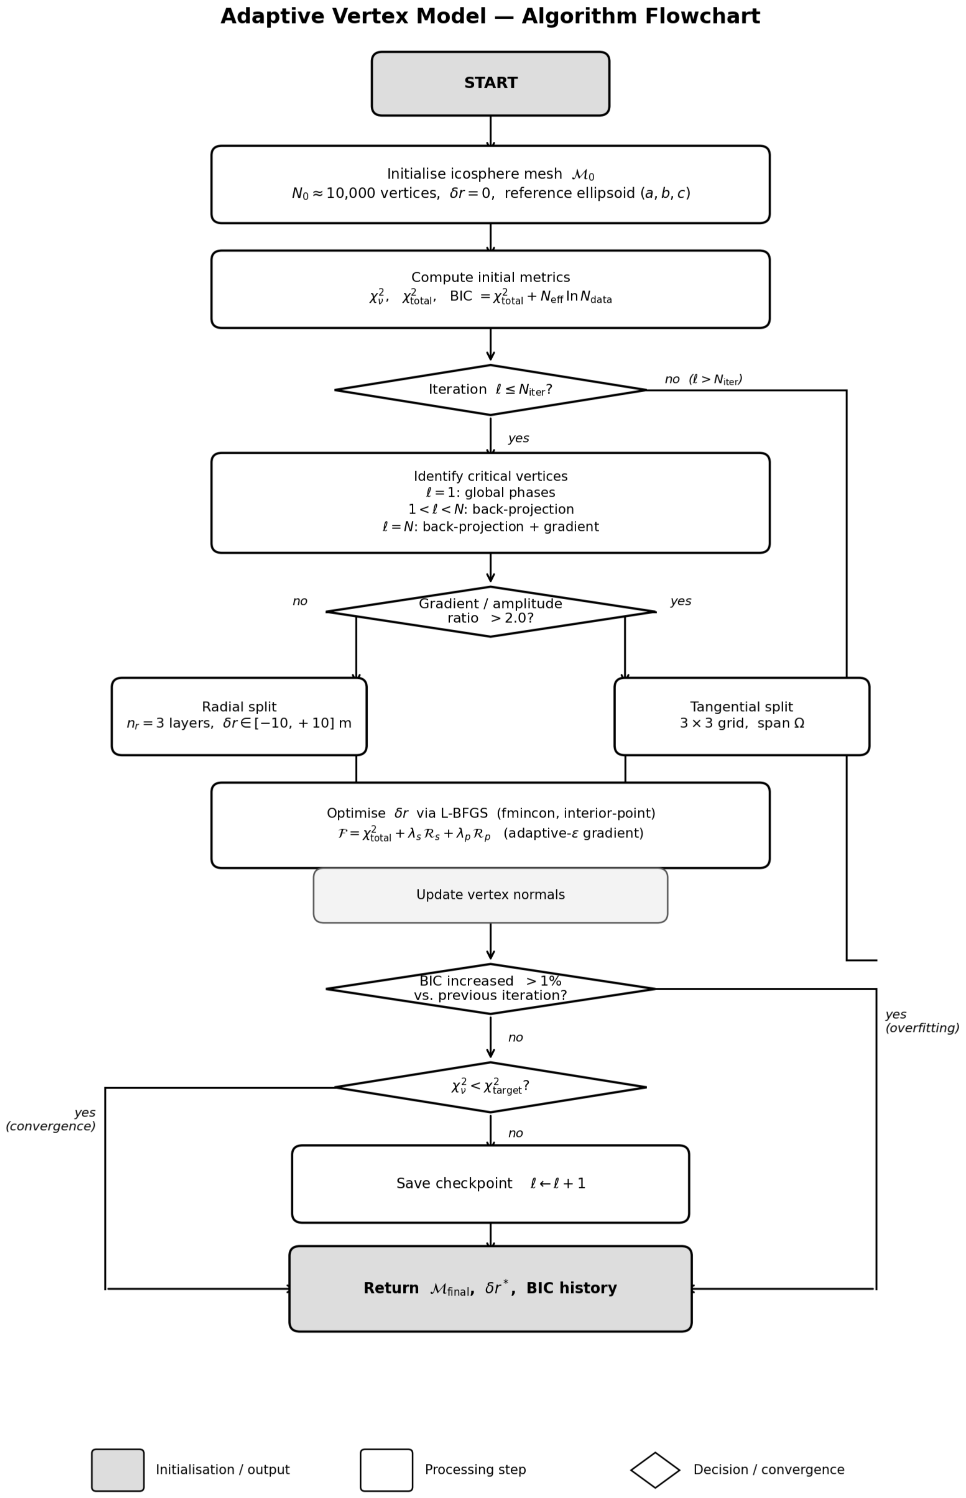

Figure 4.25: Flowchart of the adaptive vertex algorithm

prescribed target N_0. Then, the unit sphere is mapped to the reference ellipsoid by coordinate scaling $(x, y, z) \to (ax, by, cz)$. Vertex normals are initialized analytically from the ellipsoid implicit function. All perturbation are set to zero, so the starting model coincides with the reference ellipsoid.

After each vertex splitting operation and after each optimization step, the vertex normals are recomputed from the current mesh geometry using area-weighted averaging of adjacent face normals. For vertex i belonging to the face set $\mathcal{F}_i$:

$$\hat{n}i = \frac{\mathbf{N}i}{||\mathbf{N}i||}, \quad \text{with} \quad \mathbf{N}i = \sum_{f \in \mathcal{F}_i} (\mathbf{e}_{f,1} \times \mathbf{e}_{f,2}) \tag{4.41}$$

where $\mathbf{e}_{f,1}$ and $\mathbf{e}_{f,2}$ are edge vectors of face f emanating from a common vertex. The result is normalized to unit length and its sign verified for outward consistency against the radial direction of the corresponding ellipsoid base point. This ensures that the Lambertian scattering law is always evaluated with normals consistent with the actual perturbed geometry.

At each stage of the algorithm, the relation between the perturbed mesh and the observed spectra is quantified by the cost function:

$$\mathcal{J}(\delta\mathbf{r}) = \chi^2_\nu + \lambda_s \mathcal{R}s + \lambda_p \mathcal{R}_p \tag{4.42}$$

where χ^2_ν is given by Equation **4.32**. The regularization terms $\mathcal{R}_s$ and $\mathcal{R}_p$ control the solution space. The former is a smoothness penalty term, based on the normalized graph Laplacian of the mesh: it penalizes spatially incoherent perturbation patterns:

$$\mathcal{R}_s = \frac{1}{2} ||\mathbf{L}\,\delta\mathbf{r}||^2 \tag{4.43}$$

while the latter is a soft proximity constraint that penalizes displacements exceeding $r_{\max} = 25$ m:

$$\mathcal{R}_p = \frac{1}{2} \sum_i \max(0, |\delta r_i| - r_{\max})^2 \tag{4.44}$$

The weights λ_s and λ_p rule the trade-off between data fidelity and surface smoothness and are selected by $k = 5$ fold cross-validation on the rotational phases before the refinement loop. The N_{spec} spectra are divided into five folds by uniform phase distribution, ensuring proportional phase coverage in each fold. For each candidate pair (λ_s, λ_p), the mesh is reinitialized to the smooth ellipsoid, optimized on the training fold, and the resulting validation χ^2_ν is recorded. The pair minimizing the mean validation error across folds is adopted for the subsequent refinement loop.

The gradient of the cost function $\mathcal{J}$ with respect to δr_i is computed by forward finite differences, with an adaptive step that ensures that the gradient is estimated at a scale comparable to the local perturbation amplitude. The cost function is minimized using the L-BFGS algorithm [215], which approximates the inverse Hessian matrix from the history of gradient differences, enabling adaptive step sizes without manual tuning of a learning rate.

Two geometrically distinct splitting operations are available and selected adaptively:

- Radial splitting (Figure **4.26**) replaces vertex $\mathbf{v}_0$ with $n_r = 3$ new vertices displaced along the local normal at uniform offsets spanning $[-10, +10]$ m. The adjacent triangles are removed and replaced by a fan triangulation connecting the new vertices to the ordered neighbor ring. Radial splitting introduces local radial degrees of freedom without altering the angular coverage of the mesh, and is applied to vertices where the dominant residual is a spectral amplitude discrepancy;

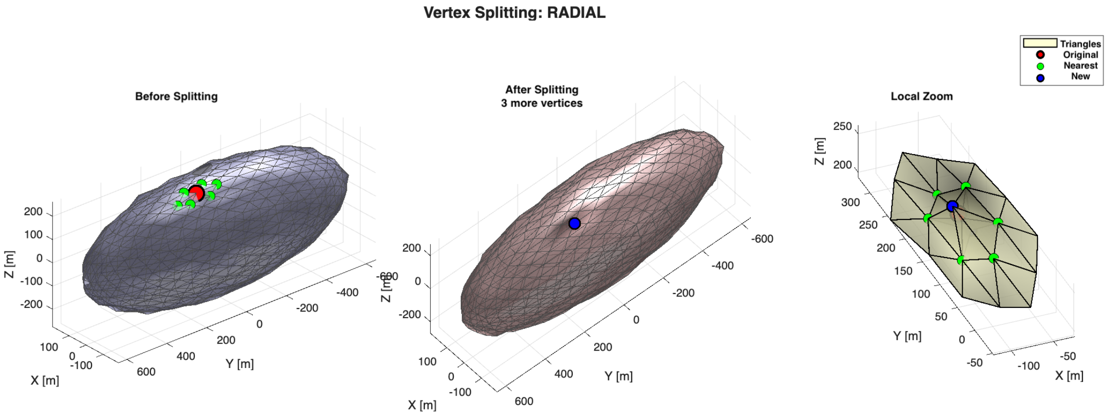

Figure 4.26: Sketch of the radial splitting procedure

- Tangential splitting (Figure **4.27**) replaces $\mathbf{v}_0$ with a 3×3 grid of new vertices distributed over an angular region of half-width $\Omega/2 = 0.75 l_0/r_0$ centered on (θ_0, ϕ_0), where l_0 is the mean local edge length and r_0 is the ellipsoid radius. Each new vertex is placed on the reference ellipsoid at its updated coordinates and inherits the radial perturbation of the original vertex. Tangential splitting increases local angular resolution and is applied where the residual reflects a spectral shape error rather than a purely amplitude error.

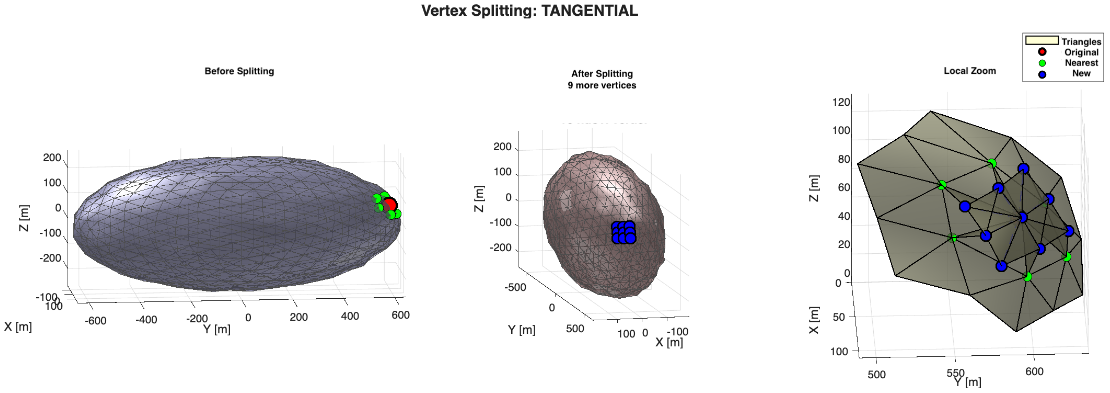

Figure 4.27: Sketch of the tangential splitting procedure

After each split, vertex normals are updated, degenerate faces are removed, and the mesh topology is validated.

To identify the critical vertices, three identification methods are applied adaptively at different stages of the algorithm:

- the global phase residual method (Method A) computes the norm of spectral residuals for each observed rotational phase and maps phases in the top quartile of the residual distribution to the corresponding visible surface hemisphere. This method is computationally inexpensive and is used only in the first iteration.

- The spectral backprojection method (Method B) evaluates the Doppler frequency f_D of each visible vertex for every rotational phase. The spectral residual at the frequency bin closest to f_D is assigned as the local residual of vertex i at phase k. A mean amplitude residual $\bar{R}_i$ is then computed over all phases in which the vertex was visible.

- The spectral gradient analysis (Method C) accumulates residuals on the smoothed spectral derivative $\partial S/\partial f$, identifying vertices responsible for shape errors in the spectrum rather than amplitude mismatches. This method is used only in the final iteration.

The first iteration uses Method A with radial splits only. Iterations 2 to $N_{\text{iter}} - 1$ use Method B, where the split type (radial or tangential) is determined by the local ratio of gradient residuals to amplitude residual: a ratio exceeding 1.0 triggers a tangential split, otherwise a radial split is applied. The final iteration combines Methods B and C. After all splits in each iteration, $\delta\mathbf{r}$ is optimized on the updated mesh and vertex normals are recomputed.

Two stopping criteria are evaluated at the end of each iteration. The first is the convergence of χ^2_ν below a prescribed target. The second is the Bayesian Information Criterion [216]:

$$\text{BIC} = \chi^2 + N_{\text{eff}} \cdot \ln\left(N_{\text{data}}\right) \tag{4.45}$$

if the BIC increases by more than 1% between consecutive iterations, the newly introduced vertices are assumed to be absorbing spectral noise rather than recovering physical shape features, and the refinement loop is terminated. The BIC criterion mitigates but does not eliminate overfitting, particularly if the spectral noise exhibits systematic structure from residual RFI or instrumental artifacts that the model can absorb with spatially localized vertex perturbations. Hence, the improvement in χ^2_ν should therefore be interpreted with caution, since a fraction of it may reflect fitting of structured noise rather than recovery of true surface topography. The backprojection localization strategy and the cross-validated regularization framework represent the methodological contributions of this modeling level and are directly transferable to datasets with higher information content, such as delay-Doppler images or multi-aspect CW observations.

4.5.10 Adaptive vertex model results

The adaptive vertex model described in Section **4.5.9** represents the most flexible level of the shape reconstruction hierarchy. Unlike the uniform patch parameterizations of the radial and vectorial perturbation models, the adaptive vertex model represents the asteroid surface as a triangulated mesh whose resolution is not fixed a priori but grows iteratively in the regions where the spectral residuals are largest, concentrating degrees of freedom where the data carry genuine shape information.

The initial mesh contains approximately 1000 vertices uniformly distributed on the reference ellipsoid via the icosphere method; after ten iterations of vertex splitting and L-BFGS optimization, the final mesh contains $N = 2682$ vertices and has undergone 60 individual vertex-splitting operations.

The spectral fits for the same six representative rotational phases used in the previous sections are shown in Figure **4.28**. The improvement in fit quality with respect to all preceding model levels is substantial and systematic: the reduced chi-squared decreases to $\chi^2_\nu = 1.08$, the closest approach to unity achieved across the entire reconstruction hierarchy. The residual excess can be explained by spectral fluctuations on scales smaller than the angular resolution of the mesh. It therefore cannot be absorbed by any finite-resolution shape model. The three-dimensional shape recovered by the adaptive vertex model is shown in Figure **4.29**. The perturbation field ranges from $\delta r_{\text{min}} = -25.0$ m to $\delta r_{\text{max}} = 4.9$ m, with a standard deviation $\sigma_{\delta r} = 8.8$ m.

The most prominent morphological feature of the recovered shape is the strong dominance of negative perturbations across virtually the entire visible surface. The color map is almost entirely

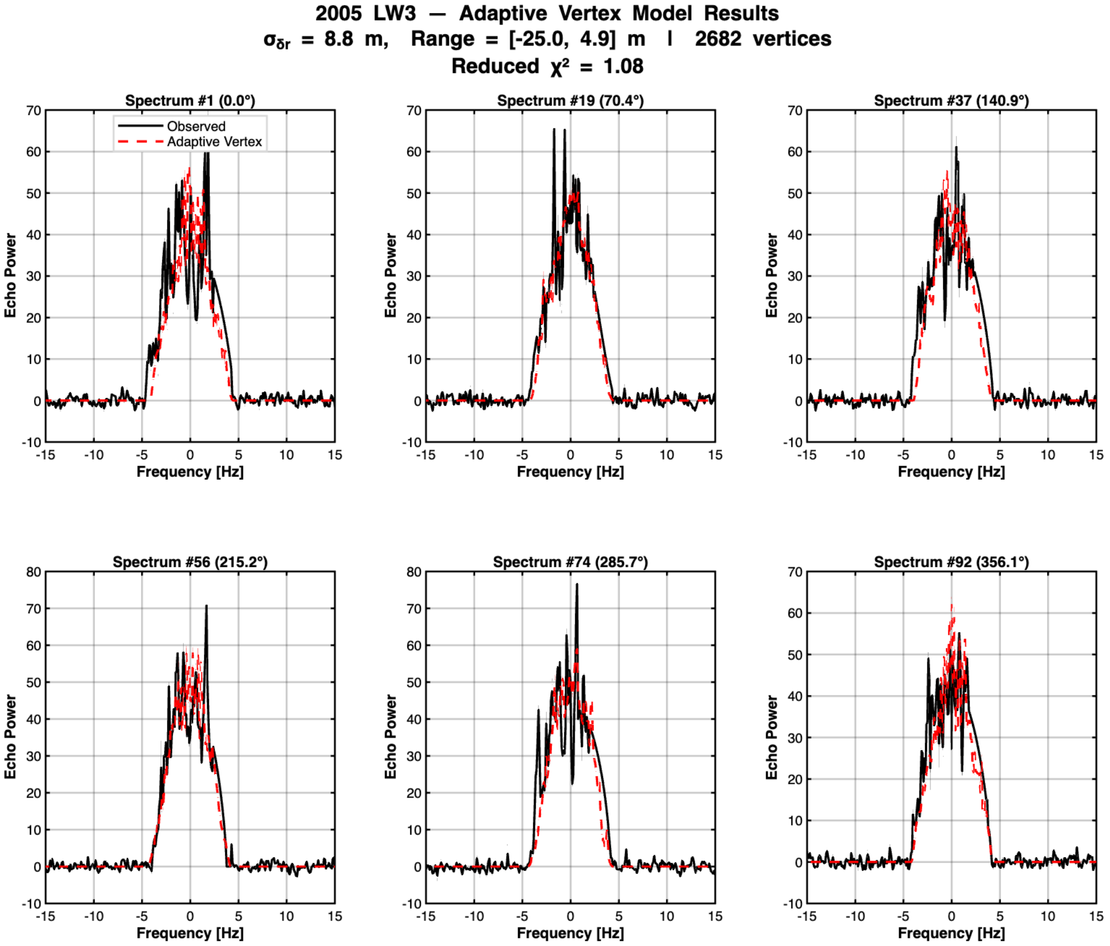

Figure 4.28: Simulated adaptive vertex model spectra (dotted red line) vs real data (black line). Frequency resolution: 0.1 Hz, integration time: 10 minutes.

blue in all six views, with only isolated white patches concentrated near the equatorial belt in the side view (YZ plane) and near mid-latitudes in the front view (XZ plane), and there is no coherent region of positive perturbation anywhere on the body. This systematic negative bias indicates that the true surface of 2005 LW3 is, on average, more compact than the reference ellipsoid in the radial direction: the smooth ellipsoid, whose semi-axes were fixed from the convex hull reconstruction and the single-parameter c optimization, slightly overestimates the effective scattering volume of the asteroid as seen in the CW echo. This is physically consistent with the presence of concave surface regions, such as craters or flat facets, that reduce the backscattered power relative to a smooth convex body of the same overall dimensions.

The spatial distribution of the perturbation field exhibits a coherent large-scale pattern rather than a noise-like texture, which supports the interpretation that it reflects genuine surface topography. The polar region visible in the side and front views shows the deepest negative perturbations. The equatorial band, which is the most strongly constrained region in a CW observation, displays a heterogeneous mix of near-zero and moderately negative perturbations with isolated white patches that break the azimuthal symmetry of the reference ellipsoid, suggesting the presence of surface features at the decametric scale. As noted throughout this analysis, the high-latitude perturbations are substantially less constrained by the CW data and should be interpreted with caution; the BIC stopping criterion provides a partial safeguard against overfitting but cannot fully discriminate between genuine polar topography and structured noise absorbed by the locally refined mesh.

Comparing these results with those from the radial perturbation model of Section **4.5.6**, several differences are evident and they are due to the inherent geometric configuration of the two approaches.

First, let's consider the polar regions. The radial perturbation model identifies a concentration of positive perturbations near the north pole as the most prominent departure from the reference ellipsoid. The adaptive vertex model, by contrast, places its deepest negative perturbations in the same polar region. The most likely cause is a discretization artifact inherent to the uniform 8×16 patch grid. At the resolution of that parameterization, the entire polar cap is covered by a single patch spanning approximately $20°$ in latitude. This patch accumulates the spectral residuals projected from every rotational phase in which the polar region is visible, and the SLSQP optimizer displaces it outwards because it is the only degree of freedom available to reduce the chi-squared in that region. When the same data are fitted with the adaptive mesh, the polar region is covered by many vertices whose individual perturbations are allowed to vary independently under a smoothness regularization: the optimizer distributes the same spectral information across a finer spatial grid, and the solution converges on a pattern of localized negative perturbations rather than a single positive bump.

At equatorial latitudes, where the CW spectra provide the strongest and most direct constraints through the Doppler edge positions and the shape of the spectral envelope at every rotational phase, the two models are broadly consistent: both recover a perturbation pattern of moderate amplitude with no dominant sign, corresponding to a surface that departs from the reference ellipsoid by comparable amounts in both the positive and negative radial directions. Hence, the divergence between the two models is confined almost entirely to the polar regions, where the CW observables carry minimal geometric information about the out-of-plane shape. This spatial coincidence between the zone of model disagreement and the zone of minimal data constraint is not accidental: it constitutes strong internal evidence that the polar discrepancy is a reconstruction artifact rather than a reflection of genuine surface topography, and it confirms that the shape information recoverable from this dataset is robustly captured at low latitudes by both models regardless of their different parameterizations and regularization strategies.

A second, physically more significant factor is the behavior of the reference ellipsoid itself. The polar semi-axis $c = 226$ m was determined by optimizing a single free parameter against the full spectral dataset using a smooth Lambertian scattering model. If the true asteroid is slightly more compact than this ellipsoid in nearly all directions, as the dominantly negative perturbation field of the adaptive model suggests, then the 128-patch model lacks the degrees of freedom to express a global volumetric offset efficiently: the optimizer distributes the mismatch unevenly across patches, some of which absorb it as a positive perturbation. The adaptive model, with 2682 independent vertices and a globally consistent regularization, recovers the solution closest to the true shape given the constraints of the data, and that solution is one in which the reference ellipsoid slightly overestimates the effective scattering volume of 2005 LW3 in most surface regions.

The increase in the standard deviation of the perturbation field from $\sigma_{\delta r} = 4.8$ m to $\sigma_{\delta r} = 8.8$ m is consistent with this interpretation. The adaptive mesh resolves surface variations at a spatial scale significantly finer than the patch grid, so it captures local topographic relief that the coarser model averaged out within each patch. The larger $\sigma_{\delta r}$ is therefore a more faithful estimate of the true RMS surface roughness at the resolution of the final mesh.

In summary, the two models are not inconsistent: they describe the same surface at different spatial resolutions and with different regularization constraints. The polar positive feature of the radial perturbation model is most plausibly a resolution artifact, and the globally negative field of the adaptive vertex model is physically the more reliable result, indicating that the reference ellipsoid provides a slight overestimate of the asteroid's true radial extent. The improvement in χ^2_ν confirms that the adaptive model describes the data more faithfully, and the spatial coherence of the recovered perturbation field supports the interpretation that it reflects genuine surface topography rather than numerical noise.

To conclude, the adaptive vertex model provides the highest-fidelity reconstruction of the three-dimensional shape of 2005 LW3 achievable from the available CW radar dataset. The final χ^2_ν indicates that the model describes the data at nearly the noise level, and the recovered perturbation field is dominated by a coherent pattern of small surface concavities.

Together with the results of the simpler model levels, this reconstruction establishes a self-consistent, progressively refined picture of 2005 LW3 as a nearly equidimensional body with mild equatorial elongation, and surface topographic relief at spatial scales of several tens of meters.

4.6 VLBI observations of NEOs

4.6.1 VLBI technique fundamentals

Radio interferometry is a high-resolution observational technique that leverages the principle of wave interference to overcome the diffraction-limited angular resolution of individual antennas, effectively synthesizing a virtual aperture through the combination of signals from an array of sensors. Aperture synthesis interferometry, specifically, provides angular resolution proportional to $\lambda/D_{\max}$, where $D_{\max}$ is the maximum projected separation (baseline) between receiving stations, rather than building a physically continuous aperture of diameter $D_{\max}$. The technique records the complex voltage at each antenna and computes, in post-processing, the cross-correlation between every pair of stations. The theoretical foundation of this method lies in the Van Cittert-Zernike

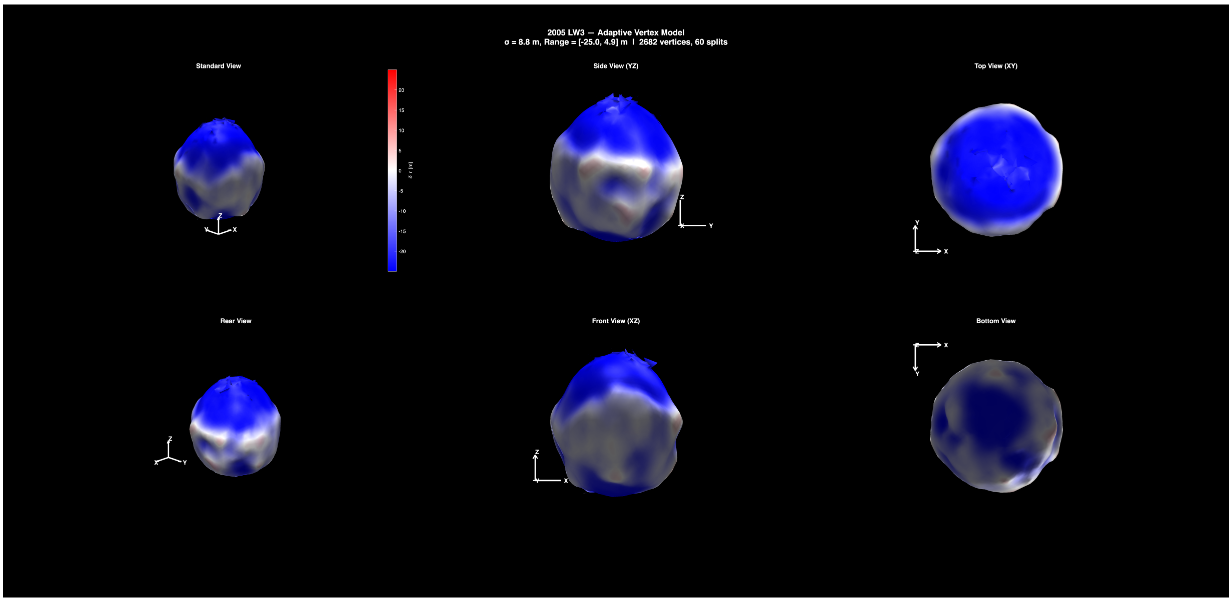

Figure 4.29: Three-dimensional shape of the adaptive vertex model output. The surface is color-coded according to the radial perturbation δr at each vertex relative to the reference ellipsoid: blue indicates regions where the reconstructed surface lies below the reference ellipsoid, corresponding to depressions, concavities, or flat facets; white corresponds to near-zero perturbation, i.e. surface elements that locally coincide with the smooth ellipsoid, and red indicates regions where the surface extends beyond the reference ellipsoid, corresponding to ridges, bulges, or elongated lobes.

theorem [217]. In the far-field approximation (Fraunhofer region), the theorem states that the mutual coherence function of a distant, incoherent source is the Fourier transform of its intensity distribution. In the context of radio astronomy, this implies that each pair of antennas acts as a two-element interferometer sensitive to a specific spatial frequency of the sky brightness distribution $I(\boldsymbol{\sigma})$, determined by the projected baseline b_{ij}. Assembling measurements from many baselines populates the Fourier plane (the uv-plane) and allows reconstruction of $I(\sigma)$ up to the diffraction limit of the longest baseline [218].

The fundamental observable is the visibility, V_{ij}, defined as the cross-correlation of the voltage streams from two stations i and j:

$$V_{ij} \propto \int_S G_{ij}(\boldsymbol{\sigma}) I(\boldsymbol{\sigma}) \exp\left(-j\frac{2\pi\nu}{c}(\mathbf{b}_{ij}\cdot\boldsymbol{\sigma})\right) d\Omega \tag{4.46}$$

where ν is the observation frequency and G_{ij} is the product of the normalized antenna beam patterns (primary beams), $\mathbf{b}_{ij}$ the baseline vector separating the two stations and $\boldsymbol{\sigma}$ is the unit vector pointing toward the source position in the sky. Each visibility thus corresponds to one Fourier component of $I(\boldsymbol{\sigma})$ at a specific spatial frequency. For compact sources, like asteroids, the beam pattern may be treated as constant across the source extent and factored out of the integral, so that the visibilities measure the Fourier transform of the sky brightness directly [219].

One of the most capable operational VLBI networks is the Very Long Baseline Array (VLBA) (Figure **4.30**), operated by the National Radio Astronomy Observatory (NRAO). It comprises ten 25-m paraboloid antennas distributed across North America and the Pacific Ocean, with maximum baselines of approximately 8600 km, and operates over a frequency range from 300 MHz to 90 GHz [220]. Image reconstruction from visibility data is performed in two stages. A dirty map is first formed as the inverse Fourier transform of the uv-plane populated with the measured

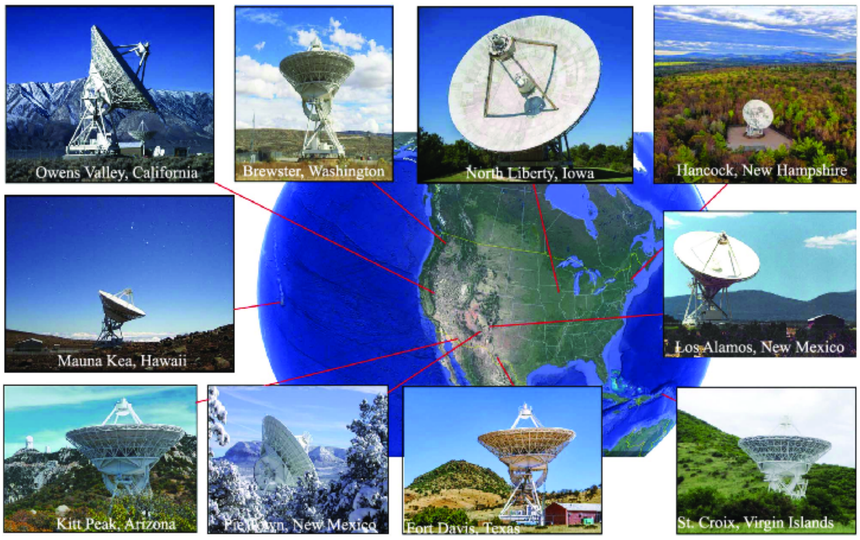

Figure 4.30: Very long baseline array (VLBA) network. From [221].

visibilities and zeros elsewhere. For computational efficiency, the visibilities are gridded to a uniform Fourier grid and the FFT is applied [222]. The dirty map is the convolution of the true brightness distribution with the synthesized beam (point spread function), whose sidelobe structure reflects incomplete uv-coverage of the interferometer. Deconvolution is typically performed via the CLEAN algorithm [223], which iteratively identifies and removes the dominant point-source contributions, and allows for estimating $I(\boldsymbol{\sigma})$, provided that signal-to-noise per synthesized beam is sufficient to distinguish real astronomical features from noise and artifacts. When antenna gain errors are non-negligible, self-calibration can further correct residual amplitude and phase errors, at the cost of losing absolute positional and flux information [224].

A non-trivial complication for Solar System targets is that the standard formulation of aperture synthesis assumes planar incoming wavefronts, i.e. a source in the far field. If the observed object is at a distance d that does not satisfy Equation **1.9**, significant phase errors arise from wavefront curvature. These errors must be corrected by applying near-field geometric delays to each baseline, effectively refocusing the array at the known target distance. For the VLBA at 2380 MHz, this boundary lies at approximately 8000 AU [225], far beyond any Solar System object. Provided that the asteroid ephemeris is accurate to track the signal phase (typically requiring a timing precision at the sub-microsecond level), these corrections restore coherent correlation between the signals across the array [222].

4.6.2 Application to Asteroid Radar Observation

Although asteroids are not intrinsically bright radio emitters, they can become intense radar sources when illuminated by high-power ground-based transmitters. The exploitation of interferometric techniques in NEO radar observations creates an observing geometry orthogonal to conventional methods, such as CW or delay-Doppler imaging. While delay-Doppler resolves the target along the line-of-sight (range) and along the Doppler-projected direction, radar-VLBI probes the plane of the sky at right angles to both [225]. The combination of these techniques, in principle, provides full three-dimensional positional information about scattering centers on the asteroid surface.

For a bistatic configuration in which the transmitter illuminates the target and a VLBI array records the reflected echo, three simultaneous constraints must be satisfied for radar interferometric

imaging.
The first condition is that the target must be angularly resolved by the array:

$$\frac{2R}{d} \geq \frac{\lambda}{D_{\max}} \tag{4.47}$$

where R is the effective target radius, d is its geocentric distance and λ the signal wavelength. If the target is smaller than one synthesized beam, no spatial structure is recovered. However, even in this unresolved case, the observation remains scientifically valuable as it allows for high-precision astrometry of the echo centroid, providing critical constraints on the asteroid's orbital parameters.
The second condition is that the target must not be over-resolved. The number of independent visibility measurements for an observation with N stations is $N(N-1)/2$. After self-calibration, which removes $2N$ phase and amplitude degrees of freedom, approximately $N(N-1)-2N$ degrees of freedom remain available in total. A reliable image reconstruction typically requires the number of independent model parameters (e.g., brightness elements or "pixels" in the image) to remain below roughly half this number. If the radar echo is not resolved in Doppler frequency, the interferometric non-over-resolution condition is:

$$\frac{2R}{d} \leq \frac{N\lambda}{2D_{\max}} \tag{4.48}$$

For the VLBA augmented by the Green Bank Telescope ($N = 11$), this limits the recoverable complexity to approximately 44 independent brightness elements on the target. Since each synthesized beam requires one such element, the target angular diameter must subtend fewer than about 6 beams in the unresolved-spectral case. If the echo is additionally resolved in Doppler frequency, each spectral channel corresponds to a narrow strip on the target in which the degrees of freedom per is approximately $N(N-1)/2$. For Doppler resolved echoes, the interferometric non-over-resolution condition then relaxes to:

$$\frac{2R}{d} \leq \frac{(N^2-3N)\lambda}{4D_{\max}} \tag{4.49}$$

allowing imaging of a broader class of targets, albeit at the expense of distributing the echo power across many spectral bins and thus requiring a stronger target to reach a sufficient signal-to-noise ratio in each frequency channel [225].
The third condition is that the asteroid's expected radar echo flux density, derived from the bistatic radar equation (Equation **1.68**), must strictly exceed the minimum detectable threshold of rms noise level of the synthesized image.
These three conditions together leave only a small fraction of known NEOs as viable candidates. To provide a concrete example, an assessment for targets approaching Earth between 2009 and 2012 identified at most twelve objects suitable for Arecibo-transmit/ VLBA-receive observations and only two for Goldstone transmit, assuming a radar albedo of 0.1 and a rotation period of 4 hours [225]. The strong distance dependence of echo power means that favorable geometries are rare and typically brief, occurring near close-approach epochs.
Even when the aforementioned conditions are satisfied, a fundamental physical obstacle prevents plane-of-sky imaging of resolved targets: the radar speckle phenomenon. Each point on the asteroid surface reflects the incident radar wave with a phase determined by its position and surface roughness. The coherent superposition of contributions from all illuminated points produces a

spatially random interference pattern (the speckle field) at the Earth. The characteristic spatial scale of this pattern is:

$$L_{\text{speckle}} = D_{\text{target}} \frac{\lambda}{d} \tag{4.50}$$

where D_{target} is the physical diameter of the asteroid. Stations separated by a baseline $b > L_{\text{speckle}}$ receive mutually uncorrelated signals: while one station may record a strong echo, another station at a different location may simultaneously record near-zero power [225]. For a resolved target, the sky brightness distribution $I(\boldsymbol{\sigma})$ perceived by each station is different, violating the spatial incoherence assumption underlying aperture synthesis. Averaging over many independent speckle realizations to recover the true $I(\boldsymbol{\sigma})$ would require integration times far exceeding the echo coherence time, making self-calibration impractical.

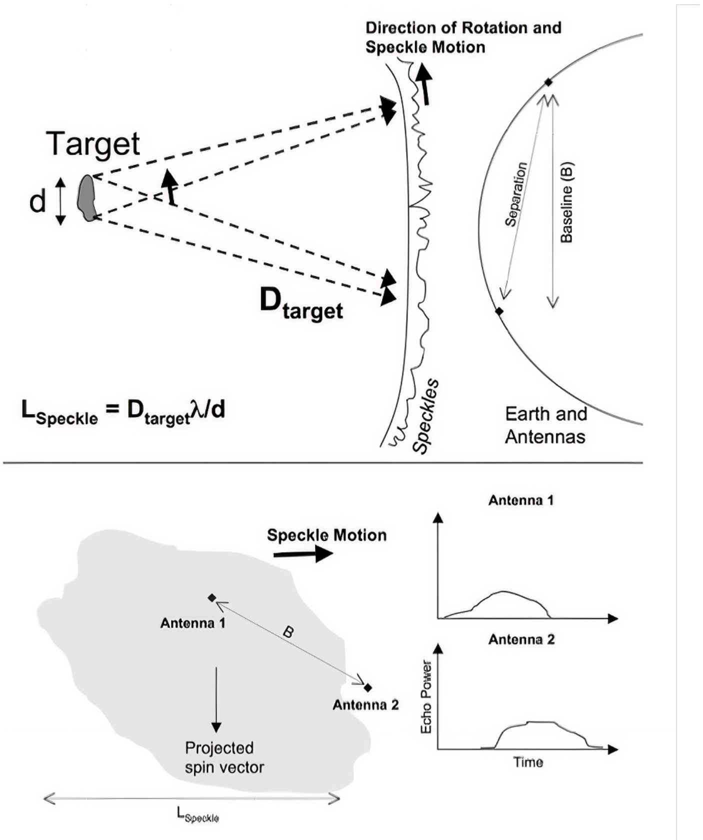

Figure 4.31: **Top**: The reflected radio wave from each point on the target's surface forms a wavefront and these interfere constructively or destructively, producing the random pattern of bright and dark speckles at the Earth. As the asteroid rotates, the phase at each point on the surface changes, moving the speckles in the same direction as the surface. **Bottom**: A single bright speckle passing over two antennas with baseline $b < L_{\text{speckle}}$ and the echo power at each as a function of time. From [225].

The speckle scale can, in principle, be increased by resolving the target in Doppler frequency, which confines each spectral channel to a narrow isodoppler strip and thereby enhances L_{speckle} in the along-strip direction. However, this approach simultaneously requires both a favorable array geometry and echo power strong enough to sustain detection across many spectral bins, a combination that is rarely achievable in practice.

4.6.3 Existing observations and State of the Art

Apart from the theoretical appeal of the technique, experimental demonstrations of radar-VLBI for Solar System objects remain rare, and the existing literature is dominated by a series of experiments coordinated by the Keldysh Institute of Applied Mathematics in Moscow. This research group operated an international network of receiving stations distributed across East Europe and Italy, using the RT-70 planetary radar at Evpatoria (Ukraine) as ground-based transmitter.

Radar VLBI asteroids was first proposed by [226], who proposed combining monostatic radar ranging with VLBI correlation to simultaneously measure range, radial velocity, angular position, and angular rate of near-Earth objects. The technique was subsequently applied to a series VLBI radar (VLBR) campaigns of space debris, asteroid close approaches and inner planets starting in 2001 [183, 227, 228], demonstrating that post-processing of radar VLBI data could reduce the uncertainty in orbital parameters by up to an order of magnitude relative to optical-only solutions, and that rotation periods and characteristic sizes of rotating space debris could be estimated from the quasi-periodic modulation of the interferometer response amplitude.

The most detailed and directly relevant demonstration is the experiment performed during the close approach of asteroid 2012 DA14 on 2013 February 15–16 [229]. At the time of minimum distance (≈ 33000 km), the RT-70 Evpatoria transmitter illuminated the target at 5 GHz frequency (C-band) with a power of ~ 70 kW, in both monochromatic (CW) and linear frequency-modulated (LFM) modes. Reflected echoes were recorded by the RT-32 antennas at Medicina (Italy) and Irbene (Latvia), with baselines of approximately $1600 - 1700$ km. The data were processed in three sequential stages: first, autocorrelation spectra at each station to characterize the echo; second, bistatic cross-correlation between transmitted and received signals to measure the Doppler shift; and finally full VLBI cross-correlation between the two receiving stations. All three stages yielded detections across the entire observation campaign, which extended from the minimum approach to a distance ten times greater. However, the continuous observation windows were shorter than the estimated rotation period (> 8 h), preventing an independent estimation of the object's size.

A companion experiment in April 2012, targeting nine space debris objects with the same RT-70 Evpatoria transmitter and a four-station receiving network (Irbene, Medicina, Urumqi, Simeiz), demonstrated systematic measurements of Doppler frequency shifts consistent with independently known orbital parameters and confirmed the viability of LFM signals for precise delay measurement and angular coordinate determination [228].

Taken together, these works establish the observational feasibility of radar-VLBI for Solar System and Earth orbiting objects, but they also expose its current limitations: the experiments were conducted with a small network of moderate-aperture antennas ($25 - 32$ m), baselines of at most a few thousand kilometers, and the limited transmitter power (as in our scenario). Thus, the achievable angular resolution and sensitivity in this configuration are substantially below what would be required for morphological characterization of most NEOs; the primary scientific product remains the astrometric refinement of orbital solutions rather than surface imaging.

4.7 VLBI proposal for 2025 FA22 observation

4.7.1 Target selection and scientific objectives

The asteroid 2025 FA22 was discovered on 29 March 2025 by the Pan-STARSS 2 survey. With an absolute magnitude of 21.5, it has an estimated diameter of approximately 160 m (within a factor of two), placing it firmly in the category of PHA. Its close approach to Earth on 18 September 2025, at a distance of 0.0056 AU (approximately 2.2 lunar distances), represented the nearest passage of a NEA larger than 100 m since 2022, and made it by a substantial margin the strongest radar target of the year, with an expected signal-to-noise ratio of order 120000 per track on the Goldstone planetary radar schedule [230], [231]. At the time, no prior physical characterization of the object was available: its low optical observability in the months before the encounter meant that no rotation period, shape, or spin state information could be obtained before the radar campaign. The close approach window was therefore the only practical opportunity to collect this data.

The scientific objectives of the proposed CW-VLBI radar campaign were the following. The first was the precise astrometric localization of the radar echo centroid in the plane of the sky. The interferometric angular resolution achievable with the proposed network is the order of one-hundredth of an arcsecond, approximately one order of magnitude better than the best high-precision ground-based optical astrometry of asteroids, which is generally limited to ~ 0.1 arcseconds. This precision directly refines the orbital solution and trajectory prediction. Critically, it also breaks the 180° pole ambiguity intrinsic to monostatic delay-Doppler imaging: because the Doppler spectrum alone cannot distinguish north from south on the asteroid, the plane-of-sky localization provided by VLBI is required to unambiguously reconstruct the spin vector and to correctly orient any three-dimensional shape model derived from radar data.

The second objective was the spectroscopic analysis of the reflected CW signal, yielding accurate Doppler frequency shifts per receiving antenna. These measurements provide the line-of-sight velocity of the asteroid with high precision and, through temporal analysis of the echo bandwidth, constrain the rotation rate and surface velocity distribution (Section **2.2.1**).

The third objective was speckle-based spin characterization. As described in Section **4.6.2**, the speckle pattern produced by a rotating asteroid has a characteristic spatial scale defined by Equation **4.50**. At the close approach distance of 2025 FA22, this scale is approximately 220 km [225], which is well matched to the shortest baselines of the proposed receiving network. By measuring the temporal and spatial statistics of the speckle intensity fluctuations across antenna pairs, it is possible to determine the rotation period, spin axis orientation, and sense of rotation of the target.

A fourth, programmatic objective was methodological: to demonstrate and refine the end-to-end technical strategies for EVN-based radar observations of near-Earth objects, contributing to the development of a future coordinated European planetary radar capability in the framework of the study described in Section **4.3**.

4.7.2 Observation Planning and Network Configuration

The observation was designed around a multistatic geometry in which the transmitter was the 70-m DSS-63 dish (Madrid). The antenna transmitted a CW signal at 7167 MHz fixed frequency with 20 kW of power in RCP, during a dedicated transmission window from 06:40 to 09:10 UT on 18 September 2025. It should be noted that this frequency belongs to the X-band according to the standard DSN nomenclature. However, in the radio astronomy community, it is conventionally categorized within the C-band (specifically C-high). To ensure signal reception, the receiving network was composed of a specific subset of the European VLBI Network (EVN) and the e-MERLIN array, selected based on their receiver bandwidth compatibility with the 7.2 GHz radar echo. The geographic distribution of the transmitting and receiving stations is shown in Figure

4.32, which illustrates the multistatic configuration of the obsercation network across Europe. The choice of CW rather than chirped waveforms was necessitated by the constraints of the DSS-63 transmitter. Unlike the DSS-14 (Goldstone) in GSSR configuration, the Madrid station is not equipped to transmit specific modulated signals for radar observations (e.g., BPC or Chirp). However, this limitation did not adversely affect the campaign: CW concentrates all transmitted power in a single spectral line, maximizing the SNR ratio on each baseline for interferometric detection. While the spatial information is absent in a single-dish CW observation, it can be recovered through the interferometric measurements.

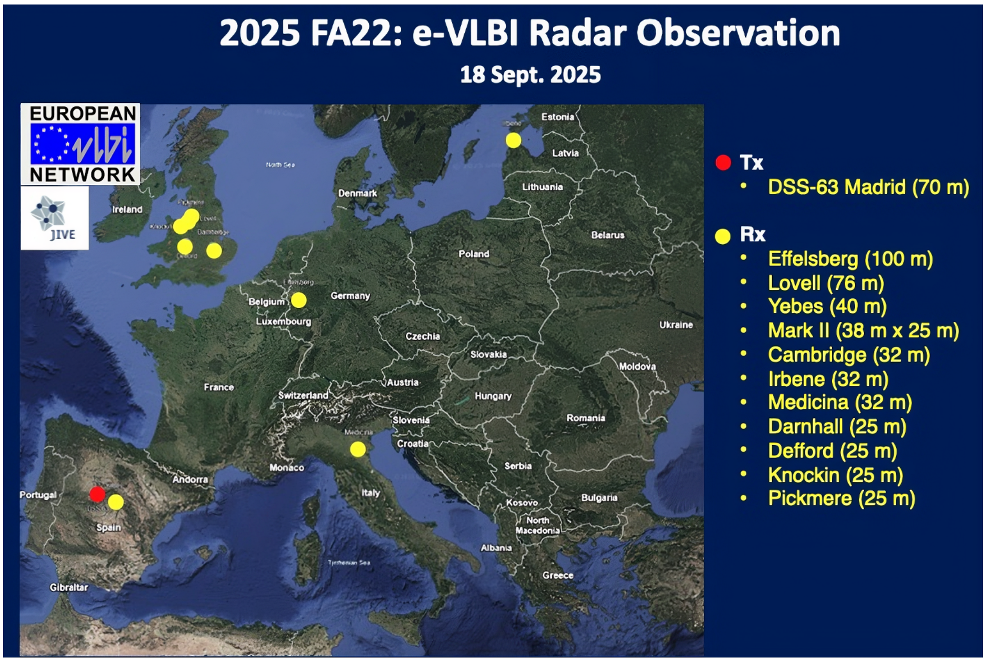

Figure 4.32: Network configuration for the e-VLBI radar observation of asteroid 2025 FA22 on 18 September 2025. The transmitter (red dot) is the DSS-63 in Madrid, which transmitted a CW signal at 7167 MHz with 20 kW of power. The receiving network (yellow dots) consists of a subset of the European VLBI Network (EVN) and the e-MERLIN array, including Effelsberg (100 m), Lovell (76 m), Yebes (40 m), Mark II (38 m × 25 m), Cambridge (32 m), Irbene (32 m), Medicina (32 m), Darnhall (25 m), Defford (25 m), Knockin (25 m), and Pickmere (25 m).

A Target of Opportunity (ToO) proposal was submitted to the EVN Program Committee and approved with high priority. The receiving network combined elements from the European VLBI Network (EVN) and the e-MERLIN array. This dual-array configuration was specifically designed to simultaneously access two complementary observational regimes. On the long EVN baselines (spanning up to several thousand kilometers), the angular resolution is sufficient to potentially resolve the target and image individual surface scattering regions, constrain shape asymmetries, and localize the rotation pole. On the short e-MERLIN baselines (ranging from tens to a few hundred kilometers), the target is unresolved. However, its echo centroid can be tracked with high astrometric precision, and the speckle scale of ~ 220 km is accessible for spin state characterization. This combination of EVN and e-MERLIN in a single radar campaign was without precedent in the literature, providing a multi-scale view of the target's physical and dynamical properties.

The full network configuration is summarized in Table **4.9**. For each station, the table reports the antenna diameter, the frequency range of the C-high receiver, and the System Equivalent Flux Density (SEFD). The SEFD is a fundamental measure of a radio telescope's sensitivity, representing the flux density of a source that would produce a response equal to the system noise. It is defined as SEFD $= 2k_B T_{\text{sys}}/A_{\text{eff}}$, where k_B is the Boltzmann constant, T_{sys} is the system temperature and A_{eff} is the effective collecting area. A lower SEFD indicates a more sensitive station.

The total requested observing window extended from 06:10 to 09:40 UT, including time for fringe-finder calibrations before and after the DSS-63 transmission. The common visibility window of all selected receiving antennas during the DSS-63 transmission is shown in Figure **4.33**, which displays the azimuth and elevation trajectories of 2025 FA22 for each station during the radar transmission. Data acquisition used the DBBC backend with four baseband channels of 16 MHz bandwidth in dual polarization at 2-bit quantization, for a total aggregate bit rate of 2 Gbps per station, compatible with real-time e-VLBI transfer to the JIVE correlator in Dwingeloo (the Netherlands) for the calculation of the interferometric visibilities. In addition, the observation proposal specifically requested that raw VDIF data streams from individual antennas be stored for independent off-line Doppler and speckle analysis.

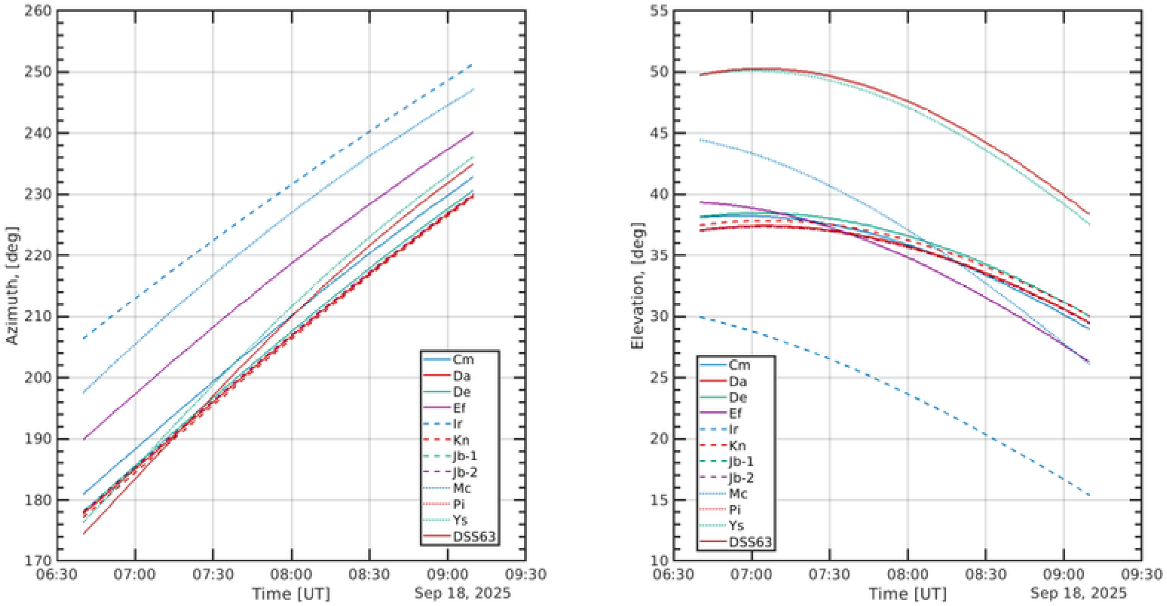

Figure 4.33: Pointing coordinates (azimuth and elevation) of asteroid 2025 FA22 for the selected receiving antennas during the radar transmission from the DSS-63 antenna on 18 September 2025 (06:40–09:10 UT)

Off-line correlation will be performed using the PRIDE (Planetary Radio Interferometry and Doppler Experiments) software pipeline developed at JIVE [233], originally designed for radio science experiments involving spacecraft in the Solar System, requiring near-field wavefront curvature corrections. The expected echo flux density, computed assuming the lower bound on the asteroid diameter as a conservative estimate, was approximately 125 Jy concentrated in the narrow spectral line containing the CW echo, yielding a minimum SNR of 15 for the shorter baselines, even under conservative assumptions.

Table 4.9: Receiving network for the 2025 FA22 VLBI radar campaign. SEFD values at C-band are taken from the EVN Status Tables [232]. The C-high frequency range and agility data refer to the 5 cm band (6.0/6.7 GHz), which covers the DSS-63 transmission frequency of 7167 MHz. The parameter SCHED indicates that frequency switching is controllable directly from the observing schedule without operator intervention.

Station	**Code**	**Array**	**Diam.** (m)	**SEFD** (Jy)	**C-high Rx range** (MHz)	**Switch time**	**SCHED** ctrl.	**Notes**
EVN stations								
Effelsberg	Ef	EVN	100	20	4000–9300	40 s	yes	linear pol.[a]
Lovell	Jb-1	EVN/e-MERLIN	76	70	4000–7500	180 s	yes	$\lambda \geq 6$ cm only
Jodrell Mk II	Jb-2	EVN	38×25	300	4000–7500	180 s	yes	
Yebes	Ys	EVN	40	180	4500–9000	<5 min	yes	C-X rx since 2022
Medicina	Mc	EVN	32	170	5900–7000	70 s[b]	yes	
Irbene	Ir	EVN	32	430	4500–8800	—	—	4 selectable 1 GHz bands
e-MERLIN out-stations								
Cambridge	Cm	e-MERLIN	32	175	4000–7500	60 min	no	
Darnhall	Da	e-MERLIN	25	325	4000–7500	60 min	no	
Defford	De	e-MERLIN	25	1000	4000–7500	60 min	no	mesh dish[c]
Knockin	Kn	e-MERLIN	25	325	4000–7500	60 min	no	
Pickmere	Pi	e-MERLIN	25	325	4000–7500	60 min	no	
Transmitter (bistatic Tx only)								
Robledo DSS-63	Rob70	-	70	—	7167 MHz CW	—	—	20 kW, RCP

[a] Effelsberg C-high band receiver operates in linear polarization; The correlation output is converted to circular polarization before data distribution.

[b] Switch time of 240 s applies for transitions between prime and secondary focus (i.e. to/from 6 cm wavelength); within secondary-focus bands the switch time is 70 s and is SCHED-controlled.

[c] Defford is a mesh dish and operates at reduced sensitivity at C-band relative to the other e-MERLIN out-stations.

4.7.3 Technical challenges and implementation

The observation presented a series of non-standard technical challenges, each of which required specific solutions and collectively these revealed the current limitations of the VLBI infrastructure originally designed for distant astrophysical sources.

At a geocentric distance of 0.0056 AU at closest approach ($\sim$ 838000 km), the parallax displacement of 2025 FA22 between geographically separated antennas is non-negligible. Stations separated by thousands of kilometers observe the target at measurably different apparent sky positions simultaneously. Consequently, a single set of pointing coordinates, as used in standard astrophysical VLBI observations, was therefore inapplicable. The magnitude of this parallax effect is quantified in Figure **4.34**, which displays the differences in equatorial coordinates of 2025 FA22 as seen from each receiving antenna relative to Effelsberg. This station was arbitrarily selected as the common reference point for this comparison, as it represents the largest and most sensitive receiving antenna in the network. Each antenna required an individually computed ephemeris-based pointing sequence, generated by a radar observation planning software based on the NAIF-SPICE routines using trajectory kernels downloaded from JPL Horizons. Furthermore, the target's proper motion during the 2.5 hour observation window reached approximately 8700 arcsec/h in right ascension and 4400 arcsec/h in declination, a rate that sweeps several degrees across the sky during a single session and makes standard sidereal tracking entirely inadequate. The right ascension and declination of the target, together with their rates of change as seen from Effelsberg during the

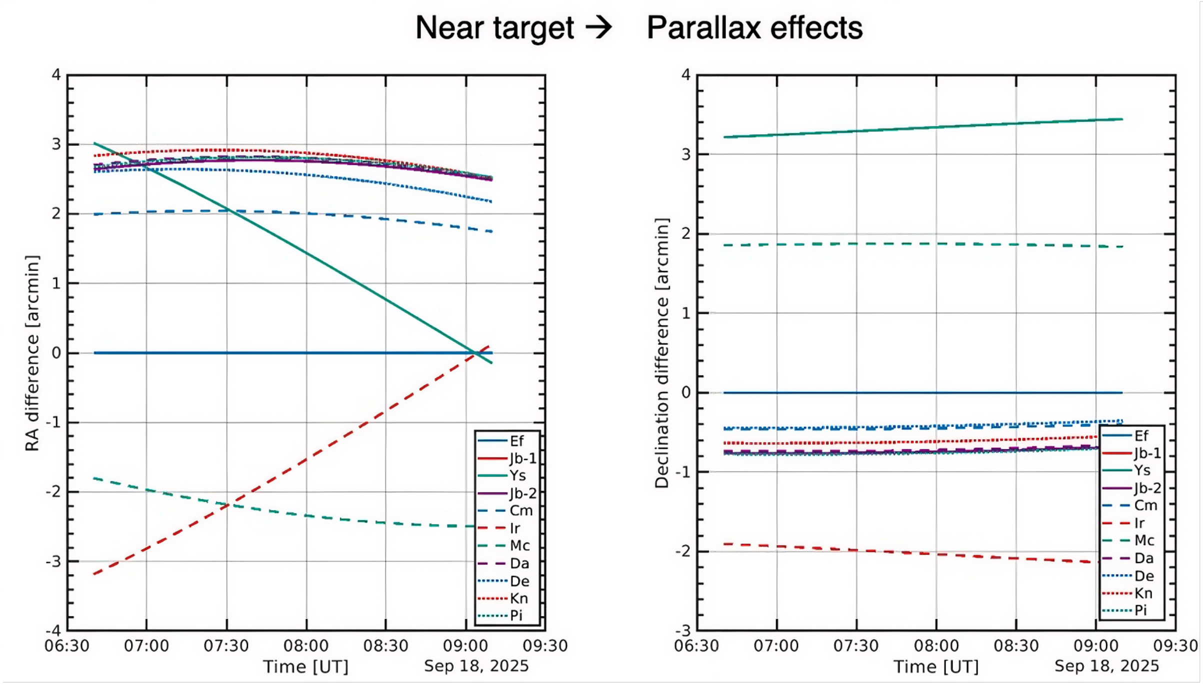

Figure 4.34: Differences in equatorial coordinates of 2025 FA22 as seen from each receiving antenna, relative to Effelsberg, during the observation window on 18 September 2025. The parallax effect arising from the target's proximity requires each station to use an individually computed pointing ephemeris.

transmission window, are shown in Figure **4.35**, illustrating the non-sidereal proper motion that precluded standard tracking.

The practical implementation of differential non-sidereal tracking varied considerably across the network, as the existing VLBI Field System did not provide a standardized mechanism for this mode of operation. Effelsberg and Lovell bypassed the Field System entirely, feeding pointing coordinates from JPL Horizons directly to the antenna control systems. The e-MERLIN array used sequences of azimuth-elevation commands computed externally. Yebes implemented an interpolated table of Az-El-time positions from JPL Horizons, with $\sim 10-15$ seconds of extra time required at each target-to-calibrator switch. Medicina implemented a *radecrate* command within the Field System snap file, providing continuous tracking with explicit rate input from JPL Horizons or SPICE ephemerides. This solution proved to be particularly robust for this type of observation. Irbene used a timed RA-Dec sequence in the snap file.

Pointing accuracy depended critically on the quality of the available orbital solution. As shown in Figure **4.36** (left panel), JPL solution #24, distributed in July 2025, proved insufficiently precise for reliable pointing of the larger, more directional antennas: the 1-sigma uncertainty ellipse of the asteroid's position was comparable with the Effelsberg antenna beam. The updated JPL solution #25, incorporating additional astrometric observations, was released on 16 September 2025, less than 48 hours before the start of the observation. As illustrated in the right panel of Figure **4.36**, this new solution drastically reduced the positional uncertainty, bringing the error ellipse well within the antenna beam and thereby guaranteeing successful tracking. However, this required a short-notice update to the observing schedule, which had already been distributed to all stations on 15 September. This event underlines a structural challenge for NEO radar campaigns: some

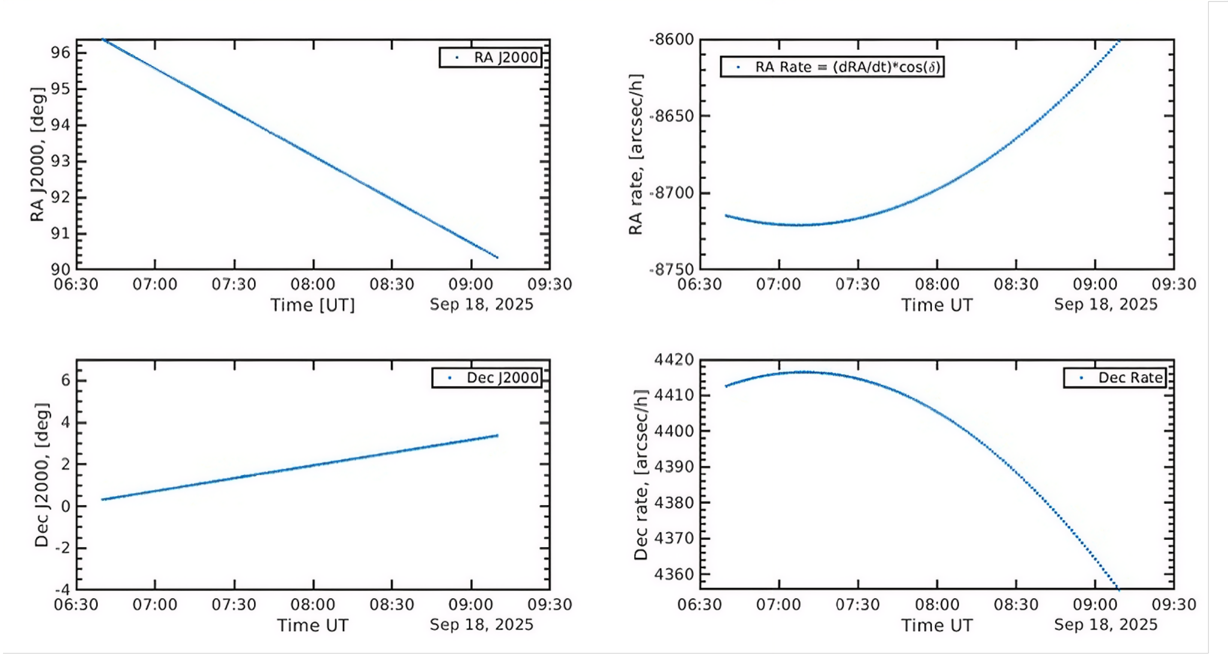

Figure 4.35: The J2000 Right Ascension (RA) and Declination (Dec), and their respective rates of change for asteroid 2025 FA22. The panels show RA (**top left**) and its rate corrected for declination, $(d\mathrm{RA}/dt) \cdot \cos(\mathrm{Dec})$ (**top right**), Declination (**bottom left**) and its rate, dDec/dt (**bottom right**). All values are for the Effelsberg radio telescope during the DSS-63 transmission window (06:40–09:10 UT, 18 September 2025).

close-approach targets are intrinsically poorly constrained orbitally until close to the encounter date, precisely when the observation is imminent and operational flexibility is lowest.

Due to the fast proper motion of the target, a single phase reference calibrator could not serve the entire observation. As the target moved, the angular offset between target and calibrator grew, degrading the phase transfer. To maintain the required angular proximity, the observation relied on switching calibrators along the target's trajectory during the phase-referenced part of the observation.

The observing schedule itself was carefully structured into distinct phases to accommodate both continuous target tracking and phase-referenced interferometry, while mitigating the challenges posed by the asteroid's fast proper motion. Before the radar transmission, the session began with prolonged observations of two fringe-finder calibrators (4C39.25 and 3C84) to synchronize the stations and establish absolute delay references. This was immediately followed by initial scans of the two selected phase-reference calibrators (J0613-0008 and J0604+0204), compact background sources with flux densities above 100 mJy, located within 2.5 degrees of the target track. During the first part of the DSS-63 transmission, the network executed three long, contiguous 30-minute scans on the asteroid without interspersed phase calibrators. This continuous tracking strategy was specifically designed to optimize single-dish measurements and speckle tracking, which benefit from uninterrupted scans.

In the second part of the transmission window, the strategy switched to an interferometric phase-referenced mode. The adopted strategy alternated target scans with calibrator observations in a 5-minute cycle (3 minutes on target, 2 minutes on calibrator), consistent with the C-band atmospheric coherence time of 4-6 minutes. Due to the asteroid's rapid motion across the sky, maintaining the

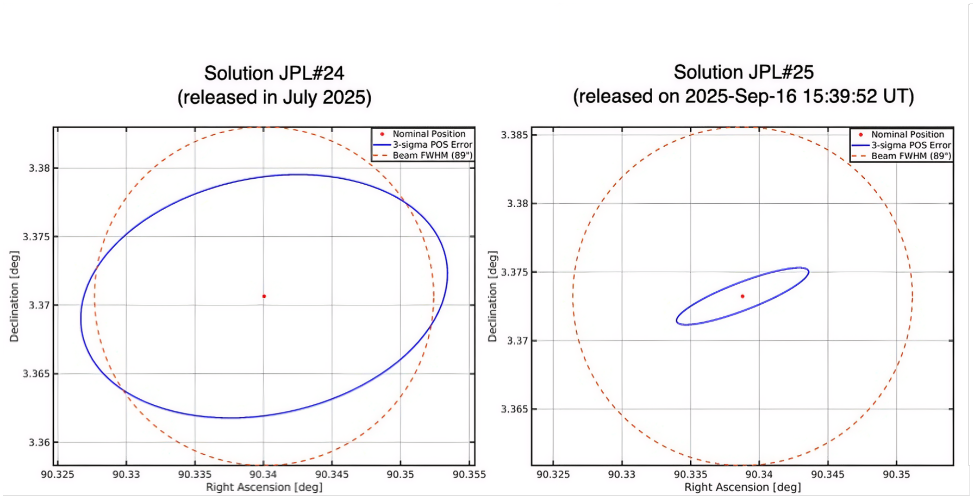

Solution JPL#24
(released in July 2025)

Solution JPL#25
(released on 2025-Sep-16 15:39:52 UT)

Figure 4.36: Comparison of maximum pointing uncertainties for 2025 FA22 during the radar observation window. **Left panel:** The 1-σ target position uncertainty ellipse (solid blue line) from JPL solution #24 is comparable with the Effelsberg antenna beam (dashed red line) centered on the nominal pointing (red cross). **Right panel:** The updated JPL solution #25 drastically reduced the uncertainty ellipse to well within the antenna beam, ensuring successful pointing.

required angular proximity (within 2.5°) required a calibrator change. Consequently, the first phase calibrator (J0613-0008) was replaced by the second (J0604+0204) as the asteroid moved along its trajectory. After the radar transmission concluded, the network performed alternating scans of the two phase calibrators alone, and the session closed with a final fringe-finder observation (3C84). Post-observation coordination with JPL revealed that the DSS-63 transmission had to be interrupted during two brief intervals for satellite avoidance purposes, resulting in the actual transmission being divided into three sub-windows: 06:40–07:20 UT, 07:27–07:36 UT, and 07:48–09:10 UT, for a net illumination time of 131 minutes (approximately 2.2 hours). Furthermore, the effective radar echo reception time was even shorter than this net illumination time. During the phase-referenced portion of the observation, the receiving antennas spent a significant fraction of the time pointing away from the asteroid to observe the phase calibrators. The complete observational timeline is summarized in Figure **4.37**.

4.7.4 Observation outcome

The observation was executed on 18 September 2025 as planned. At present, the echo detection has been confirmed at the Lovell telescope, where the echo flux density of ~ 125 Jy, concentrated in a spectral interval of a few Hz around the Doppler-shifted CW line, was detected with high SNR. A representative spectrum of this detection is shown in Figure **4.38**.

In practice, the majority of participating antennas encountered difficulties in implementing the differential tracking requirements within their existing Field System configurations, resulting in tracking degradation or data losses for portions of the session or even for the entire observation. Unfortunately, these losses included both the e-MERLIN array, which precluded the use of the short baselines for speckle tracking, and the Effelsberg antenna, which was the most sensitive receiving element in the network.

The correctly-tracking antennas - whose data are currently at JIVE for correlation and near-field reduction using the PRIDE pipeline - are expected to support at minimum an astrometric VLBI

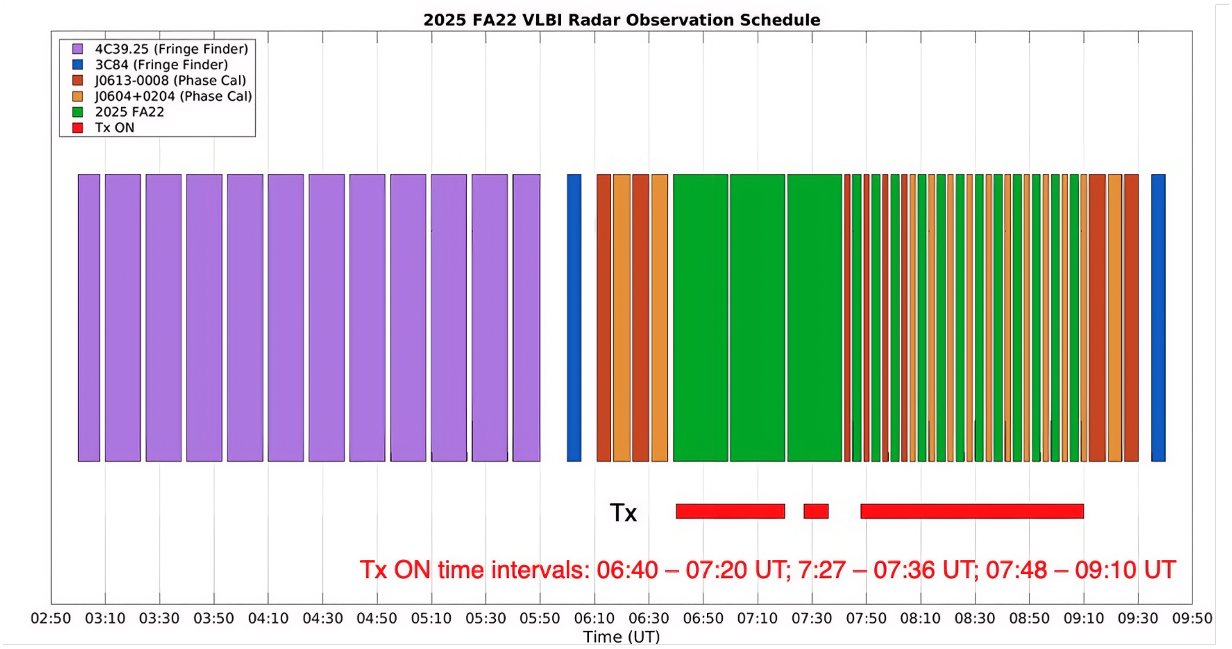

Figure 4.37: 2025 FA22 VLBI Radar Observation Schedule

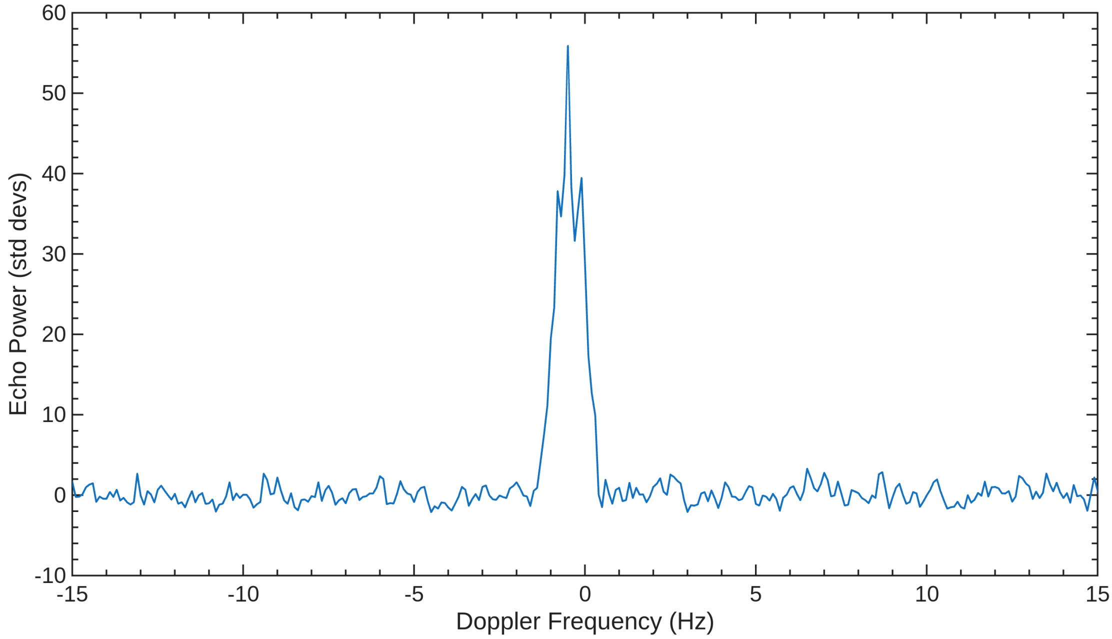

Figure 4.38: Integrated power spectrum at 0.1 Hz frequency resolution, 4 min integration time, of 2025 FA22 echo recorded at Lovell on 18 September 2025 at 06:53:48 UT. Echo power on the vertical axis is plotted in standard deviations of the background noise. The Doppler frequency on the horizontal axis is relative to the ephemeris-based estimate frequency of the echo from the asteroid COM.

solution, i.e. a plane-of-sky position of the echo centroid at one or more epochs, and Doppler spectral measurements enabling rotation analysis. Whether the SNR and uv-coverage from the well-performing baselines will support a resolved interferometric image of surface scattering regions, will be determined by the correlation results currently in progress.

The lessons learned from the campaign could be directly implemented in future observations, and are summarized here as operational recommendations.

First, the most critical infrastructure need is the implementation of a standardized differential tracking mode for non-sidereal targets within the VLBI Field System, since its absence was the primary source of tracking failures across multiple stations and required heterogeneous station-specific workarounds.

Second, the experience with the late orbital solution update suggests that future campaigns should establish, in advance, a formal protocol for rapid schedule updates and coordinate distribution shortly before the observation.

These technical improvements, combined with the scheduling strategies, calibration scheme, and network configuration validated by this campaign, constitute a concrete and tested roadmap toward a European capability for VLBI radar characterization of near-Earth asteroids.

Chapter 5

Conclusions

This thesis has developed and validated advanced radar signal processing methodologies across two complementary domains of planetary science: subsurface sounding of Mars through orbital radar data, three-dimensional characterization of near-Earth asteroids through ground-based CW observations. Despite their apparent diversity, the theoretical and methodological foundation is the same: extract physical information from radar measurements in regimes of limited data, high noise, and ill-posed inversion. Together, this work constitutes a contribution to the broader goal of planetary radar as a tool for both scientific investigation and planetary defense.

5.1 MARSIS subsurface sounding

The detection of a bright basal reflector beneath the Martian south polar ice cap is, at the moment of writing, the first known evidence of liquid water on present-day Mars. The investigation of the Ultimi Scopuli region of the Martian south pole began with a critical re-examination of the existing MARSIS ground-processing chain.

MARSIS transmits a frequency-modulated chirp signal in the HF band, operating across four sub-bands centered at $1.8, 3.0, 4.0$, and 5.0 MHz with an instantaneous bandwidth of 1 MHz each, yielding a free-space range resolution of approximately 150 m per sub-band, which reduces to roughly $50 - 100$ m within the low-permittivity polar ice. The dataset used in this thesis consist in the raw MARSIS echoes, transmitted to ground, with the full range compression performed in the ground processing chain. This choice preserves the maximum flexibility for processing.

The identification of a remnant phase term causing incoherent integration of orbital data represented the key diagnostic step that motivated a complete revision of the raw data processing pipeline. The noise sources identified in this work - demodulation artifacts, orbital radial velocity variations, and ionospheric phase delays - are not specific to the Ultimi Scopuli dataset but are inherent to the MARSIS acquisition geometry, and their characterization therefore has general validity for any MARSIS-based subsurface investigation.

The main methodological contribution of this thesis is the development and implementation of a three-dimensional frequency domain backprojection algorithm specifically optimized for the MARSIS geometry and data characteristics. Since no standard approach exists for 3D planetary radar imaging, it was necessary to survey distinct 3D reconstruction algorithms drawn from different fields - computed tomography, seismology, electron microscopy, and radio astronomy - in order to identify the most appropriate framework for MARSIS data. The theoretical foundation of the chosen approach is the formulation of the solution to the ill-posed inverse scattering problem as an approximate inverse, obtained by introducing the Born Approximation. The adjoint operator $\mathcal{A}^*$ implemented as time-domain or frequency domain backprojection requires no knowledge of the propagation medium beyond the two-way travel time from each radar position to each image point. Furthermore, among all the algorithms evaluated, it is the best suited for the irregular geometry of the Martian surface, the non-uniform along-track sampling rate characteristic of MARSIS orbits, and the altitude variations of the spacecraft during each acquisition. Residual ionospheric phase errors, provided they are small enough not to decorrelate the signal across the aperture, do not produce coherent artifacts: since each aperture position contributes independently to the 3D reconstruction, phase errors do not accumulate in a structured way.

Although the frequency-domain implementation (more flexible than the time-domain counterpart, since it doesn't require interpolation of the irregularly sampled aperture positions) constitutes the only viable choice given the characteristics of the MARSIS dataset, the validity of the Born Approximation implies that the final 3D reflectivity matrix cannot be interpreted quantitatively in absolute terms. However, given the primary goal of this work, which is the detection and spatial mapping of the bright basal reflector, this limitation is acceptable and does not affect the scientific conclusions.

To ensure that the 3D coherent integration does not amplify artifacts or noise sources, the available dataset, in a region of approximately 50 km radius around the bright reflector, was first evaluated through 2D focusing of the individual radargrams. The SNR before and after the focusing step was compared, and only orbits showing a positive ΔSNR were retained for the 3D reconstruction. In the original dataset, the L1C range compressed data, this led to just 3 orbits selectable for the reconstruction. Despite the intrinsic limitations of the tomographic reconstruction imposed by the limited angular diversity at disposal, the 3D results obtained demonstrate that the algorithm performs as designed, producing a volumetric reconstruction of the subsurface reflectivity that is inaccessible to conventional 2D focusing approaches.

The phase and noise corrections introduced in this work were then applied to the L1C data, generating a corrected dataset that responded significantly better to the 2D focusing algorithm. A new 3D volume was subsequently reconstructed using the same three orbital tracks. Comparison with the uncorrected volume reveals a consistent improvement: the corrected reconstruction is visibly more compact, with fewer but better-defined structures, consistent with improved phase coherence across the synthetic aperture. In particular, several bright spots that appeared spatially separated in the original volume are now connected, forming a single continuous bright reflector extending over several kilometers. Analysis of the depth profiles of the region provided an estimate of the apparent vertical extent of the dominant echo, found to be approximately 0.8–1 km at the -10 dB threshold. However, this value cannot be reliably interpreted as an estimate of the true geometric thickness of the reflecting layer: if the medium has a high loss tangent, as would be expected for briny water, MARSIS may lack the sensitivity to detect a return from the bottom of the interface, and the observed width would then reflect only the upper boundary of the reflector convolved with the system point spread function, rather than its full vertical extent.

The same reconstruction pipeline was subsequently applied to the full set of 27 noise- and phase-corrected orbits, substantially expanding the coverage of the synthetic aperture. The resulting volumetric reconstruction reveals again the coherent reflector at an apparent depth of approximately 3 km that is no longer confined to isolated columns above individual orbit tracks, but forms a spatially extended and geometrically connected structure elongated in the east–west direction. The geometry of this structure is qualitatively consistent with the map of anomalously high basal dielectric permittivity reported by [1] from inversion of the same orbital dataset. Also, secondary reflectors visible at slightly different depths and lateral positions in the volumetric rendering are spatially compatible with the smaller high-permittivity patches surrounding the main feature in the [1] map, suggesting the possibility that these features are not isolated bodies but rather parts of a single, laterally extended and geometrically complex basal interface.

5.2 Ground-based radar characterization of near-Earth asteroids

Ground-based radar astronomy represents one of the most powerful tools available for the physical characterization of near-Earth objects, combining the unique capability of active illumination with the geometric precision of Doppler frequency measurements. The work presented in this thesis in this domain spans three interconnected contributions: participation in the European Space Agency pilot project for NEO radar observations, the development of a complete 3D shape reconstruction hierarchy for CW-only data, and the first steps toward VLBI interferometric radar observations of asteroids.

The ESA pilot project "NEO Observation Concepts for Radar Systems" established the operational and technical framework for a European ground-based radar network capable of conducting systematic NEO observations. The contribution of this thesis to that project included participation in actual radar observation campaigns using a network of European radio telescopes as receivers, among which the 32-m "G. Grueff" antenna in Medicina. The observations of asteroids 2021 AF8 and (4660) Nereus provided concrete demonstrations of the capabilities and limitations of the European receiving network, with SNR estimates validated against real data, and established the observational baseline upon which future European NEO radar campaigns can be designed. The analysis of these datasets contributed to a broader assessment of the scientific return achievable with the current European infrastructure and of the system requirements that a dedicated radar transmitter would need to satisfy.

The case study of asteroid 2005 LW3 produced two scientifically distinct results. The first is the identification of a binary system through the detection of a secondary spectral peak in the high-resolution integrated CW spectra. The binary nature of the asteroid was not the primary objective of the observation campaign, but it emerged directly from the data processing. The detection of a satellite through a spectral signature in CW data requires both sufficient SNR and careful data pre-processing that preserves the fine structure of the echo spectrum, and its identification here validates the end-to-end data quality of the processing chain. The second result is the 3D shape reconstruction of the primary component, which constitutes the main methodological contribution of this part of the thesis.

The progressive reconstruction hierarchy developed here, from the 2D convex hull model through the rotational ellipsoid, the radial and vectorial perturbation models, and finally the adaptive vertex model, constitutes a complete and self-consistent framework for CW-only shape reconstruction, in which each level builds on the results of the previous one and the information content of the dataset is exploited in a controlled and physically motivated sequence. The convex hull and ellipsoid models establish the large-scale geometry robustly and with minimal assumptions, yielding semi-axes $a = 185$ m, $b = 191$ m and $c = 226$ m. The radial perturbation model captures surface topography at the scale of the patch grid with RMS amplitude $\sigma_{\delta r} = 4.8$ m. The vectorial perturbation model demonstrated empirically that adding unconstrained degrees of freedom actively worsen the reconstruction. This is not due to a failure in the algorithm pipeline but to an intrinsic limitation of CW data. A CW radar transmits a monochromatic, unmodulated signal and therefore provides no range resolution whatsoever: the echo from the entire visible hemisphere of the asteroid is collapsed into a single power spectrum whose only dimension is Doppler frequency. This means that the three-dimensional shape of the body must be inferred entirely from the one-dimensional distribution of echo power as a function of frequency, observed at multiple rotational phases. Several degeneracies are inherent to this configuration: surface features at the same Doppler frequency but at different ranges along the line-of-sight are indistinguishable in the echo, so the reconstruction is fundamentally insensitive to the depth distribution of scatterers. Also, the polar regions of the asteroid contribute negligibly to the Doppler spread of the echo when the subradar latitude is low, as was the case for the 2005 LW3 observations, making the out-of-plane shape poorly constrained. Third, any rotational phase interval not covered by observations corresponds to a sector of the body's silhouette that is entirely unconstrained, imposing a maximum coverage requirement of the full rotation for a reliable reconstruction. These limitations are not algorithmic but physical, and no processing technique can overcome them given CW-only data.

The adaptive vertex model, with its local refinement strategy based on spectral backprojection and BIC-controlled stopping, achieves a reduced chi-squared of $\chi^2_\nu = 1.08$, close to the noise floor of the dataset. The cross-model comparison between the radial perturbation and adaptive vertex results provided a physically interpretable picture of the reconstruction uncertainties, identifying the equatorial band as the robustly constrained region and the polar regions as inherently ambiguous, given the CW observing geometry, as discussed above.

Taken together, these results demonstrate that reliable estimation of the large-scale dimensions and overall three-dimensional figure of a sub-kilometer asteroid is achievable using only Doppler data from CW observations, provided that the observation arc covers nearly the full rotation period of the body. The identification of local concavities and surface perturbations beyond the reference ellipsoid, while physically plausible and internally consistent across model levels, must be interpreted with greater caution: the fundamental limitations of the CW observing geometry discussed above mean that fine morphological features are always reconstructed with lower confidence than the global shape parameters. A rigorous validation of this aspect of the reconstruction hierarchy will require testing the algorithm on CW datasets for asteroids that have also been observed with delay-Doppler imaging, so that the shape model recovered from Doppler-only data can be compared directly against an independent, range-resolved ground truth. This represents the most important open validation step for the methodology developed in this thesis.

5.3 VLBI observations of near-Earth asteroids

The VLBI component of this thesis represents a methodological frontier for asteroid radar astronomy. The complete feasibility study developed for the observation of asteroid 2025 FA22 includes: SNR estimates for individual telescopes and baselines, and analysis of the technical challenges specific to the VLBI-radar configuration. It establishes the observational and analytical framework for a class of experiments that has no established precedent. The near-field delay compensation, correlation, fringe detection, and phase calibration of the acquired dataset, which aim to yield one of the first interferometric radar measurements of an asteroid's angular position at milliarcsecond precision, remains the immediate next step and the subject of ongoing analysis.

A further dimension of the analysis concerns the possibility that the asteroid is spatially resolved on the intermediate and long baselines of the array. At the observing frequency and for the baseline lengths involved, the fringe visibility as a function of baseline length encodes the apparent angular size of the echo-scattering region, and differences from a point-source response would provide a direct measure of the asteroid's projected extent on the sky. In the longer term, sufficient uv-plane coverage from a future well-designed multi-station campaign could, in principle, support aperture synthesis imaging of the echo, opening a path toward morphological constraints on the scattering surface that are entirely complementary to those obtained from delay-Doppler techniques.

Finally, the framework developed here has direct applications to spin pole estimation via speckle tracking. While the aforementioned loss of the short e-MERLIN baselines precludes reaching this specific objective for the current campaign, this work therefore opens a concrete and well-defined observational option for future campaigns, and the methodology documented in this thesis provides a reusable template for planning analogous experiments on future close-approach targets that meet the stringent requirements for such observations.

5.4 Outcome

The comprehensive theoretical and algorithmic treatment developed in this thesis, spanning the full signal processing chain from raw chirp and CW waveforms to three-dimensional volumetric reconstructions, reflects the extraordinary breadth of applicability of radar as a remote sensing modality. This versatility, combined with the unique ability to penetrate surfaces, operate independently of solar illumination, and provide quantitative geometric measurements at ranges inaccessible to any passive technique, establishes radar as one of the most scientifically powerful and strategically indispensable instruments available to planetary exploration in the coming decades.

For MARSIS, the immediate priority is the extension of the 3D backprojection pipeline to the full south polar dataset, which contains hundreds of orbits over the SPLD region that have not yet been processed with the corrected and focused methodology developed here. It must be acknowledged, however, that the results achievable with the current dataset remain fundamentally bounded by the characteristics of the MARSIS instrument itself: the HF operating frequency, and the noise environment of the MARSIS orbit set hard limits on the achievable range resolution and signal-to-noise ratio that no processing algorithm, however sophisticated, can overcome. The methodology developed here extracts the maximum information content from the available data, but a qualitative improvement in subsurface characterization will require a radar sounder tailored for the purpose. Such an instrument, designed specifically for high-resolution subsurface mapping of Mars, could in principle produce volumetric reconstructions of the polar stratigraphy at metric-scale resolution across large portions of the south polar layered deposit, with the sensitivity to detect additional bright basal reflectors beyond the Ultimi Scopuli region, each of which would constitute a potential candidate for subsurface liquid water and a target of the highest priority for future Mars exploration.

For the CW shape reconstruction hierarchy, the most important open validation step is the application of the full pipeline to asteroids for which independent delay-Doppler shape models already exist. A direct comparison between the Doppler-only reconstruction and a range-resolved ground truth would quantify the fidelity of each level of the hierarchy, establish the confidence bounds on the recovered surface topography, and define the observational conditions under which the methodology can be applied reliably to targets with no delay-Doppler coverage. This validation campaign could be an additional target for future observing runs within a European NEO radar network.

For the VLBI observations, the 2025 FA22 campaign produced a concrete set of operational lessons, among which the most significant is the standardization of tracking mode for non-sidereal targets within the VLBI Field System. This improvement, combined with the scheduling strategies, calibration scheme, and network configuration validated by this campaign, constitutes a tested and immediately actionable roadmap toward a European capability for VLBI radar characterization of near-Earth asteroids.

List of Figures

List of Tables

Bibliography

[1] Sebastian Emanuel Lauro, Elena Pettinelli, Graziella Caprarelli, Luca Guallini, Angelo Pio Rossi, Elisabetta Mattei, Barbara Cosciotti, Andrea Cicchetti, Francesco Soldovieri, Marco Cartacci, Federico Di Paolo, Raffaella Noschese, and Roberto Orosei. Multiple subglacial water bodies below the south pole of mars unveiled by new marsis data. *Nature Astronomy*, 5:63–70, 2021.

[2] Mark A. Richards, James A. Scheer, and William A. Holm. *Principles of Modern Radar: Basic Principles.* SciTech Publishing, 2013.

[3] David Long and Ulaby. *F. Ulaby and D.G. Long, Microwave Radar and Radiometric Remote Sensing, University of Michigan Press, Ann Arbor, Michigan, 2014.* Michigan University, 01 2014.

[4] Fawwaz Ulaby, David Long, William Blackwell, Charles Elachi, Adrian Fung, Christopher Ruf, K. Sarabandi, Jakob Zyl, and Howard Zebker. *Microwave Radar and Radiometric Remote Sensing.* Artech House, 01 2014.

[5] E. Pettinelli, B. Cosciotti, F. Di Paolo, S. E. Lauro, E. Mattei, R. Orosei, and G. Vannaroni. Dielectric properties of jovian satellite ice analogs for subsurface radar exploration: A review. *Reviews of Geophysics*, 53(3):593–641, 2015.

[6] David J. Daniels. *Ground Penetrating Radar.* The Institution of Engineering and Technology, 2nd edition, 2004.

[7] Mark A. Richards. *Fundamentals of Radar Signal Processing.* McGraw-Hill, New York, 1 edition, jun 2005.

[8] Abhay K. Ram, George Vahala, Linda Vahala, and Min Soe. Reflection and transmission of electromagnetic pulses at a planar dielectric interface – theory and quantum lattice simulations. *AIP Advances*, 11(10):105116, 2021.

[9] David Stillman and Gary Olhoeft. Frequency and temperature dependence in electromagnetic properties of martian analog minerals. *J. Geophys. Res*, 113, 09 2008.

[10] Wlodek Kofman, Roberto Orosei, and Elena Pettinelli. Radar signal propagation and detection through ice. *Space Science Reviews*, 153:249–271, 06 2010.

[11] David E. Stillman, Joseph A. MacGregor, and Robert E. Grimm. The role of acids in electrical conduction through ice. *Journal of Geophysical Research: Earth Surface*, 118(1):1–16, jan 2013.

[12] Jérémie Mouginot, Antoine Pommerol, Pierre Beck, Wlodek Kofman, and Stephen M. Clifford. Dielectric map of the martian northern hemisphere and the nature of plain filling materials. *Geophysical Research Letters*, 39(2), 2012.

[13] Luigi Castaldo, Daniel Mège, Joanna Gurgurewicz, Roberto Orosei, and Giovanni Alberti. Global permittivity mapping of the martian surface from sharad. *Earth and Planetary Science Letters*, 462:55–65, 03 2017.

[14] Ling Zhang, Yi Xu, Renrui Liu, Ruonan Chen, Roberto Bugiolacchi, and Rui Gao. The dielectric properties of martian regolith at the tianwen-1 landing site. *Geophysical Research Letters*, 50(13):e2022GL102207, 2023. e2022GL102207 2022GL102207.

[15] Fawwaz T. Ulaby, Richard K. Moore, and Adrian K. Fung. *Microwave Remote Sensing: Active and Passive, Volume III: From Theory to Applications*. Artech House, Norwood, MA, 1986.

[16] T. P. Lester, M. L. McCall, and J. B. Tatum. Theory of Planetary Photometry. JRASC, 73:233, October 1979.

[17] P. Singh, M. Diwakar, A. Shankar, et al. A review on sar image and its despeckling. *Archives of Computational Methods in Engineering*, 28:4633–4653, December 2021.

[18] F. G. Bass and I. M. Fuks. Wave scattering from statistically rough surfaces. *Oxford Pergamon Press International Series on Natural Philosophy*, 93, January 1979.

[19] Jill Audrey Ogilvy and Harold M. Merklinger. *Theory of Wave Scattering From Random Rough Surfaces*. Philadelphia Institute of Physics, 1991.

[20] M. S. Shunmugam and V. Radhakrishnan. Selection and fitting of reference line for surface profiles. *Proceedings of the Institution of Mechanical Engineers*, 190(1):193–201, 1976.

[21] G. Rasigni, M. Rasigni, J. Palmari, C. Dussert, F. Varnier, and A. Llebaria. Statistical parameters for random and pseudorandom rough surfaces. *Journal of the Optical Society of America A*, 5(1):99–103, 1988.

[22] Francesca Ticconi, Luca Pulvirenti, and Nazzareno Pierdicca. Models for scattering from rough surfaces. In Vitaliy Zhurbenko, editor, *Electromagnetic Waves*, chapter 10. InTech, Rijeka, jun 2011.

[23] Bruce A. Campbell. *Radar Remote Sensing of Planetary Surfaces*. Cambridge University Press, Cambridge, mar 2002.

[24] Stephen O. Rice. Reflection of electromagnetic waves from slightly rough surfaces. *Communications on Pure and Applied Mathematics*, 4(2-3):351–378, 1951.

[25] Gaspar R. Valenzuela. Theories for the interaction of electromagnetic and oceanic waves — A review. *Boundary-Layer Meteorology*, 13(1-4):61–85, January 1978.

[26] A.K. Fung, Z. Li, and K.S. Chen. Backscattering from a randomly rough dielectric surface. *IEEE Transactions on Geoscience and Remote Sensing*, 30(2):356–369, 1992.

[27] David E. Smith, Maria T. Zuber, Herbert V. Frey, James B. Garvin, James W. Head, Duane O. Muhleman, Gordon H. Pettengill, Roger J. Phillips, Sean C. Solomon, H. Jay Zwally, W. Bruce Banerdt, Thomas C. Duxbury, Matthew P. Golombek, Frank G. Lemoine, Gregory A. Neumann, David D. Rowlands, Oded Aharonson, Peter G. Ford, Anton B. Ivanov, Catherine L. Johnson, Patrick J. McGovern, James B. Abshire, Robert S. Afzal, and Xiaoli Sun. Mars orbiter laser altimeter: Experiment summary after the first year of global mapping of mars. *Journal of Geophysical Research: Planets*, 106(E10):23689–23722, 2001.

[28] Mark A. Richards, James A. Scheer, and William A. Holm. *Principles of Modern Radar: Basic principles*. The Institution of Engineering and Technology, 2010.

[29] George T. Ruck, Donald E. Barrick, William D. Stuart, and Charles K. Krichbaum. *Radar Cross Section Handbook*. Springer, Durham, NC, United States, 1970.

[30] Giancorrado Brighi. *Monostatic and Bistatic Radar Techniques for Planetary Surface Exploration Across the Solar System*. Phd thesis, Università di Bologna, 2025. PhD Program in Aerospace Science and Technology, Cycle 37.

[31] Lamont V. Blake. *Radar range-performance analysis*. Artech house, 1980.

[32] Ronald N. Bracewell. *The Fourier Transform and Its Applications*. McGraw-Hill Electrical and Electronic Engineering Series. McGraw-Hill, New York, 1966.

[33] Dennis Gabor. Theory of communication. *Journal of the Institution of Electrical Engineers - Part III: Radio and Communication Engineering*, 93(26):429–457, 1946.

[34] Alan V. Oppenheim, Alan S. Willsky, and Syed H. Nawab. *Signals and Systems*. Prentice Hall, Upper Saddle River, NJ, 2 edition, 1999. General signal processing fundamentals.

[35] Ian G. Cumming and Frank H. Wong. *Digital Processing of Synthetic Aperture Radar Data: Algorithms and Implementation*. Artech House, Norwood, MA, 2005.

[36] Karol Krzempek, Piotr Jaworski, Piotr Bojęś, and Paweł Kozioł. Photothermal gas detection using a mode-locked laser signal readout. *Journal of Lightwave Technology*, 40:1–1, 07 2022.

[37] Claude E. Shannon. Communication in the presence of noise. *Proceedings of the IRE*, 37(1):10–21, 1949.

[38] B. Widrow. A study of rough amplitude quantization by means of nyquist sampling theory. *IRE Transactions on Circuit Theory*, 3(4):266–276, 1956.

[39] G. Picardi, D. Biccari, R. Seu, J. Plaut, W. T. K. Johnson, R. L. Jordan, A. Safaeinili, D. A. Gurnett, R. Huff, R. Orosei, O. Bombaci, D. Calabrese, and E. Zampolini. MARSIS: Mars Advanced Radar for Subsurface and Ionosphere Sounding. In Andrew Wilson and Agustin Chicarro, editors, *Mars Express: the Scientific Payload*, volume 1240 of *ESA Special Publication*, pages 51–69, August 2004.

[40] Petre Stoica and Randolph Moses. Spectral analysis of signals. *Prentice Hall*, 01 2005.

[41] James W. Cooley and John W. Tukey. An algorithm for the machine calculation of complex fourier series. *Mathematics of Computation*, 19(90):297–301, 1965.

[42] F.J. Harris. On the use of windows for harmonic analysis with the discrete fourier transform. *Proceedings of the IEEE*, 66(1):51–83, 1978.

[43] Merrill I. Skolnik. *Radar Handbook*. McGraw-Hill, 3 edition, 2008.

[44] Steven J. Ostro et al. Asteroid radar astronomy. In *Asteroids III*, pages 151–168. University of Arizona Press, 2002.

[45] Michael W. Busch et al. Radar Observations and the Shape of Near-Earth Asteroid 2008 EV5. *Icarus*, 212:649, 2011.

[46] Christian Wolff. Continuous wave radar. `https://www.radartutorial.eu/02.basics/Continuous%20Wave%20Radar.en.html`, 2024. Radar Tutorial. Accessed: 2026-03-13.

[47] Furuno Electric Co., Ltd. Radar basics. `https://www.furuno.com/en/technology/radar/basic/`, 2026. Accessed: 2026-03-13.

[48] Chenlu Zhang, Hua Dang, Yifeng Xiong, and Tianqi Yan. Spread spectrum algorithm resistance to wideband non-stationary interference. *The Journal of Engineering*, 2019, 10 2019.

[49] William J. Caputi. Stretch: A time-transformation technique. *IEEE Transactions on Aerospace and Electronic Systems*, AES-7(2):269–278, 1971.

[50] Donald R. Wehner. *High-Resolution Radar*. Artech House, 2 edition, 1995.

[51] Carl A. Wiley. Synthetic aperture radars—a paradigm for technology evolution. *IEEE Transactions on Aerospace and Electronic Systems*, AES-21(3):440–443, 1985.

[52] Paolo Trivero and Walter Biamino. Observing marine pollution with synthetic aperture radar. In Pasquale Imperatore and Daniele Riccio, editors, *Geoscience and Remote Sensing New Achievements*. InTech, Rijeka, feb 2010.

[53] F. Abu Bakar, J. Holmberg, T. Nieminen, et al. L, c and x band integrated synthetic aperture radar (sar) receiver. *Analog Integrated Circuits and Signal Processing*, 77(3):423–436, December 2013.

[54] Andy Harrison. *Introduction to Synthetic Aperture Radar Using Python and MATLAB*. Artech House Publishers, Norwood, MA, 2022.

[55] Richard Bamler and Philip Hartl. Synthetic aperture radar interferometry. *Inverse Problems*, 14, 01 1998.

[56] P.A. Rosen, S. Hensley, I.R. Joughin, F.K. Li, S.N. Madsen, E. Rodriguez, and R.M. Goldstein. Synthetic aperture radar interferometry. *Proceedings of the IEEE*, 88(3):333–382, 2000.

[57] J. Lee, L. Jurkevich, Piet Dewaele, Patrick Wambacq, and A. Oosterlinck. Speckle filtering of synthetic aperture radar images: A review. *Remote Sensing Reviews*, 8, 02 1994.

[58] Timo Lensu. Synthetic aperture radar — systems and signal processing: by John C. Curlander, California Institute of Technology, Jet Propulsion Laboratory, Pasadena, CA, USA and Robert N. McDonough, John Hopkins University, Applied Physics Laboratory, Laurel, MD, USA.

Publishers: John Wiley & Sons Ltd., Baffins Lane, Chichester, West Sussex PO19 1UD, United Kingdom, 1991, ISBN 0-471-85770-X. *Signal Processing*, 29(1):107–107, January 1992.

[59] D.E. Wahl, P.H. Eichel, D.C. Ghiglia, and C.V. Jakowatz. Phase gradient autofocus-a robust tool for high resolution sar phase correction. *IEEE Transactions on Aerospace and Electronic Systems*, 30(3):827–835, 1994.

[60] H.A. Zebker and J. Villasenor. Decorrelation in interferometric radar echoes. *IEEE Transactions on Geoscience and Remote Sensing*, 30(5):950–959, 1992.

[61] David J. Daniels. Ground penetrating radar. In Kai Chang, editor, *Encyclopedia of RF and Microwave Engineering*. Wiley-Interscience, Hoboken, NJ, 2005.

[62] R. H. Stolt. Migration by Fourier Transform. *Geophysics*, 43(1):23–48, February 1978.

[63] Raffaele Persico. *Introduction to Ground Penetrating Radar: Inverse Scattering and Data Processing*. Wiley-IEEE Press, 2014.

[64] Oz Yilmaz. *Seismic Data Analysis: Processing, Inversion, and Interpretation of Seismic Data*, volume 10 of *Investigations in Geophysics*. Society of Exploration Geophysicists, Tulsa, OK, 2 edition, 2001.

[65] Jen Gazdag. Wave equation migration with the phaseshift method. *Geophysics*, 43(7):1342–1351, December 1978.

[66] Caner Özdemir, Şevket Demirci, Enes Yiğit, and Betül Yilmaz. A review on migration methods in B-scan ground penetrating radar imaging. *Mathematical Problems in Engineering*, 2014:280738, 2014.

[67] Norman Bleistein. On the imaging of reflectors in the earth. *Geophysics*, 52(7):931–942, 1987.

[68] Norman Bleistein. *Mathematical Methods for Wave Phenomena*. Academic Press, Orlando, FL, USA, 1984.

[69] C.J. Leuschen and R.G. Plumb. A matched-filter-based reverse-time migration algorithm for ground-penetrating radar data. *IEEE Transactions on Geoscience and Remote Sensing*, 39(5):929–936, 2001.

[70] Damien Garcia, Louis Le Tarnec, Stéphan Muth, Emmanuel Montagnon, Jonathan Porée, and Guy Cloutier. Stolt's f-k migration for plane wave ultrasound imaging. *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control*, 60(9):1853–1867, 2013.

[71] Jon F. Claerbout. *Imaging the Earth's Interior*. Blackwell Scientific Publications, Oxford, England and Boston, 1985.

[72] C. Cafforio, C. Prati, and F. Rocca. Sar data focusing using seismic migration techniques. *IEEE Transactions on Aerospace and Electronic Systems*, 27(2):194–207, 1991.

[73] Anthony J. Devaney. *Mathematical Foundations of Imaging, Tomography and Wavefield Inversion*. Cambridge University Press, Cambridge, 2012.

[74] Mario Bertero and Patrizia Boccacci. *Introduction to Inverse Problems in Imaging.* CRC Press, Boca Raton, 1 edition, 1998.

[75] Weng Cho Chew, H. Gan, J. H. Lin, C. C. Lu, Gregory P. Otto, Jinming Song, R. L. Wagner, and William H. Weedon. Forward and inverse scattering problems in electromagnetic waves. In *Proceedings of SPIE*, volume 2524, San Diego, CA, United States, 1995. Presented at SPIE's International Symposium on Optical Science, Engineering, and Instrumentation.

[76] Akira Ishimaru. Wave propagation and scattering in random media. In *Wave propagation and scattering in random media*, 1997.

[77] Jacques Hadamard. *Lectures on Cauchy's Problem in Linear Partial Differential Equations.* Yale University Press, New Haven, 1923. Original work in English; classic reference on linear PDEs.

[78] Per Christian Hansen. Rank-deficient prewhitening with quotient svd and ulv decompositions. *BIT Numerical Mathematics*, 38(1):34–43, 1998.

[79] Per Christian Hansen. The discrete picard condition for discrete ill-posed problems. Technical report, Department of Mathematical Modelling, Technical University of Denmark, July 1989.

[80] V. A. Morozov. Choice of parameter in solving functional equations by the method of regularization. *Doklady Akademii Nauk SSSR*, 175(6):1225–1228, 1967. In Russian.

[81] Bruce E. Hansen. Testing for parameter instability in linear models. *Journal of Policy Modeling*, 14(4):517–533, 1992.

[82] Rainer Kress. Tikhonov regularization. In *Linear Integral Equations*, volume 82 of *Applied Mathematical Sciences*, pages 323–349. Springer, New York, 3 edition, 2014.

[83] Ilaria Catapano, Francesco Soldovieri, Giovanni Alli, Gaetano Mollo, and Luca Antonio Forte. On the reconstruction capabilities of beamforming and a microwave tomographic approach. *IEEE Geoscience and Remote Sensing Letters*, 12(12):2369–2373, December 2015.

[84] Gabor T. Herman. *Image Reconstruction from Projections.* Academic Press, New York, 1980.

[85] A. Micke. Natterer, f., the mathematics of computerized tomography. stuttgart, b. g. teubner and chichester etc., john wiley & sons 1986. x, 222 s., dm 72,–. isbn 3-519-02103-x and 0-471-90959-9. *ZAMM - Journal of Applied Mathematics and Mechanics / Zeitschrift für Angewandte Mathematik und Mechanik*, 67(11):580–580, 1987.

[86] Emil Wolf. Three-dimensional structure determination of semi-transparent objects from holographic data. *Optics Communications*, 1(4):153–156, 1969.

[87] Hee Shin, Ram Narayanan, and Muralidhar Rangaswamy. Ultrawideband noise radar imaging of impenetrable cylindrical objects using diffraction tomography. *International Journal of Microwave Science and Technology*, 2014, 12 2014.

[88] A. J. Devaney. A filtered backpropagation algorithm for diffraction tomography. *Ultrasonic Imaging*, 4(4):336–350, 1982.

[89] Avinash C. Kak and Malcolm Slaney. *Principles of Computerized Tomographic Imaging.* IEEE Press, New York, 1999. Electronic copy (c) 1999 by A. C. Kak and Malcolm Slaney.

[90] Tobias Owen, K. Biemann, D. R. Rushneck, J. E. Biller, D. W. Howarth, and A. L. Lafleur. The composition of the atmosphere at the surface of Mars. J. Geophys. Res., 82(B28):4635–4640, September 1977.

[91] Paul R. Mahaffy, Christopher R. Webster, Sushil K. Atreya, Heather Franz, Michael Wong, Pamela G. Conrad, Dan Harpold, John J. Jones, Laurie A. Leshin, Heidi Manning, Tobias Owen, Robert O. Pepin, Steven Squyres, Melissa Trainer, MSL Science Team, Osku Kemppinen, Nathan Bridges, Jeffrey R. Johnson, Michelle Minitti, David Cremers, James F. Bell, Lauren Edgar, Jack Farmer, Austin Godber, Meenakshi Wadhwa, Danika Wellington, Ian McEwan, Claire Newman, Mark Richardson, Antoine Charpentier, Laurent Peret, Penelope King, Jennifer Blank, Gerald Weigle, Mariek Schmidt, Shuai Li, Ralph Milliken, Kevin Robertson, Vivian Sun, Michael Baker, Christopher Edwards, Bethany Ehlmann, Kenneth Farley, Jennifer Griffes, John Grotzinger, Hayden Miller, Megan Newcombe, Cedric Pilorget, Melissa Rice, Kirsten Siebach, Katie Stack, Edward Stolper, Claude Brunet, Victoria Hipkin, Richard Léveillé, Geneviève Marchand, Pablo Sobrón Sánchez, Laurent Favot, George Cody, Andrew Steele, Lorenzo Flückiger, David Lees, Ara Nefian, Mildred Martin, Marc Gailhanou, Frances Westall, Guy Israël, Christophe Agard, Julien Baroukh, Christophe Donny, Alain Gaboriaud, Philippe Guillemot, Vivian Lafaille, Eric Lorigny, Alexis Paillet, René Pérez, Muriel Saccoccio, Charles Yana, Carlos Armiens-Aparicio, Javier Caride Rodríguez, Isaías Carrasco Blázquez, Felipe Gómez Gómez, Javier Gómez-Elvira, Sebastian Hettrich, Alain Lepinette Malvitte, Mercedes Marín Jiménez, Jesús Martínez-Frías, Javier Martín-Soler, F. Javier Martín-Torres, Antonio Molina Jurado, Luis Mora-Sotomayor, Guillermo Muñoz Caro, Sara Navarro López, Verónica Peinado-González, Jorge Pla-García, José Antonio Rodriguez Manfredi, Julio José Romeral-Planelló, Sara Alejandra Sans Fuentes, Eduardo Sebastian Martinez, Josefina Torres Redondo, Roser Urqui-O'Callaghan, María-Paz Zorzano Mier, Steve Chipera, Jean-Luc Lacour, Patrick Mauchien, Jean-Baptiste Sirven, Alberto Fairén, Alexander Hayes, Jonathan Joseph, Robert Sullivan, Peter Thomas, Audrey Dupont, Angela Lundberg, Noureddine Melikechi, Alissa Mezzacappa, Julia DeMarines, David Grinspoon, Günther Reitz, Benito Prats, Evgeny Atlaskin, Maria Genzer, Ari-Matti Harri, Harri Haukka, Henrik Kahanpää, Janne Kauhanen, Osku Kemppinen, Mark Paton, Jouni Polkko, Walter Schmidt, Tero Siili, Cécile Fabre, James Wray, Mary Beth Wilhelm, Franck Poitrasson, Kiran Patel, Stephen Gorevan, Stephen Indyk, Gale Paulsen, Sanjeev Gupta, David Bish, Juergen Schieber, Brigitte Gondet, Yves Langevin, Claude Geffroy, David Baratoux, Gilles Berger, Alain Cros, Claude d'Uston, Olivier Forni, Olivier Gasnault, Jérémie Lasue, Qiu-Mei Lee, Sylvestre Maurice, Pierre-Yves Meslin, Etienne Pallier, Yann Parot, Patrick Pinet, Susanne Schröder, Mike Toplis, Éric Lewin, Will Brunner, Ezat Heydari, Cherie Achilles, Dorothy Oehler, Brad Sutter, Michel Cabane, David Coscia, Guy Israël, Cyril Szopa, Gilles Dromart, François Robert, Violaine Sautter, Stéphane Le Mouélic, Nicolas Mangold, Marion Nachon, Arnaud Buch, Fabien Stalport, Patrice Coll, Pascaline François, François Raulin, Samuel Teinturier, James Cameron, Sam Clegg, Agnès Cousin, Dorothea DeLapp, Robert Dingler, Ryan Steele Jackson, Stephen Johnstone, Nina Lanza, Cynthia Little, Tony Nelson, Roger C. Wiens, Richard B. Williams, Andrea Jones, Laurel Kirkland, Allan Treiman, Burt Baker, Bruce Cantor, Michael Caplinger, Scott Davis, Brian Duston,

Kenneth Edgett, Donald Fay, Craig Hardgrove, David Harker, Paul Herrera, Elsa Jensen, Megan R. Kennedy, Gillian Krezoski, Daniel Krysak, Leslie Lipkaman, Michael Malin, Elaina McCartney, Sean McNair, Brian Nixon, Liliya Posiolova, Michael Ravine, Andrew Salamon, Lee Saper, Kevin Stoiber, Kimberley Supulver, Jason Van Beek, Tessa Van Beek, Robert Zimdar, Katherine Louise French, Karl Iagnemma, Kristen Miller, Roger Summons, Fred Goesmann, Walter Goetz, Stubbe Hviid, Micah Johnson, Matthew Lefavor, Eric Lyness, Elly Breves, M. Darby Dyar, Caleb Fassett, David F. Blake, Thomas Bristow, David DesMarais, Laurence Edwards, Robert Haberle, Tori Hoehler, Jeff Hollingsworth, Melinda Kahre, Leslie Keely, Christopher McKay, Mary Beth Wilhelm, Lora Bleacher, William Brinckerhoff, David Choi, Jason P. Dworkin, Jennifer Eigenbrode, Melissa Floyd, Caroline Freissinet, James Garvin, Daniel Glavin, Andrea Jones, David K. Martin, Amy McAdam, Alexander Pavlov, Eric Raaen, Michael D. Smith, Jennifer Stern, Florence Tan, Michael Meyer, Arik Posner, Mary Voytek, Robert C. Anderson, Andrew Aubrey, Luther W. Beegle, Alberto Behar, Diana Blaney, David Brinza, Fred Calef, Lance Christensen, Joy A. Crisp, Lauren DeFlores, Bethany Ehlmann, Jason Feldman, Sabrina Feldman, Gregory Flesch, Joel Hurowitz, Insoo Jun, Didier Keymeulen, Justin Maki, Michael Mischna, John Michael Morookian, Timothy Parker, Betina Pavri, Marcel Schoppers, Aaron Sengstacken, John J. Simmonds, Nicole Spanovich, Manuel de la Torre Juarez, Ashwin R. Vasavada, Albert Yen, Paul Douglas Archer, Francis Cucinotta, Douglas Ming, Richard V. Morris, Paul Niles, Elizabeth Rampe, Thomas Nolan, Martin Fisk, Leon Radziemski, Bruce Barraclough, Steve Bender, Daniel Berman, Eldar Noe Dobrea, Robert Tokar, David Vaniman, Rebecca M. E. Williams, Aileen Yingst, Kevin Lewis, Timothy Cleghorn, Wesley Huntress, Gérard Manhès, Judy Hudgins, Timothy Olson, Noel Stewart, Philippe Sarrazin, John Grant, Edward Vicenzi, Sharon A. Wilson, Mark Bullock, Bent Ehresmann, Victoria Hamilton, Donald Hassler, Joseph Peterson, Scot Rafkin, Cary Zeitlin, Fedor Fedosov, Dmitry Golovin, Natalya Karpushkina, Alexander Kozyrev, Maxim Litvak, Alexey Malakhov, Igor Mitrofanov, Maxim Mokrousov, Sergey Nikiforov, Vasily Prokhorov, Anton Sanin, Vladislav Tretyakov, Alexey Varenikov, Andrey Vostrukhin, Ruslan Kuzmin, Benton Clark, Michael Wolff, Scott McLennan, Oliver Botta, Darrell Drake, Keri Bean, Mark Lemmon, Susanne P. Schwenzer, Ryan B. Anderson, Kenneth Herkenhoff, Ella Mae Lee, Robert Sucharski, Miguel Ángel de Pablo Hernández, Juan José Blanco Ávalos, Miguel Ramos, Myung-Hee Kim, Charles Malespin, Ianik Plante, Jan-Peter Muller, Rafael Navarro-González, Ryan Ewing, William Boynton, Robert Downs, Mike Fitzgibbon, Karl Harshman, Shaunna Morrison, William Dietrich, Onno Kortmann, Marisa Palucis, Dawn Y. Sumner, Amy Williams, Günter Lugmair, Michael A. Wilson, David Rubin, Bruce Jakosky, Tonci Balic-Zunic, Jens Frydenvang, Jaqueline Kløvgaard Jensen, Kjartan Kinch, Asmus Koefoed, Morten Bo Madsen, Susan Louise Svane Stipp, Nick Boyd, John L. Campbell, Ralf Gellert, Glynis Perrett, Irina Pradler, Scott VanBommel, Samantha Jacob, Scott Rowland, Evgeny Atlaskin, Hannu Savijärvi, Eckart Boehm, Stephan Böttcher, Sönke Burmeister, Jingnan Guo, Jan Köhler, César Martín García, Reinhold Mueller-Mellin, Robert Wimmer-Schweingruber, John C. Bridges, Timothy McConnochie, Mehdi Benna, Hannah Bower, Anna Brunner, Hannah Blau, Thomas Boucher, Marco Carmosino, Harvey Elliott, Douglas Halleaux, Nilton Rennó, Beverley Elliott, John Spray, Lucy Thompson, Suzanne Gordon, Horton Newsom, Ann Ollila, Joshua Williams, Paulo Vasconcelos, Jennifer Bentz, Kenneth Nealson, Radu Popa, Linda C. Kah, Jeffrey Moersch, Christopher Tate, Mackenzie Day, Gary Kocurek, Bernard Hallet, Ronald Sletten, Raymond Francis, Emily

McCullough, Ed Cloutis, Inge Loes ten Kate, Ruslan Kuzmin, Raymond Arvidson, Abigail Fraeman, Daniel Scholes, Susan Slavney, Thomas Stein, Jennifer Ward, Jeffrey Berger, and John E. Moores. Abundance and isotopic composition of gases in the martian atmosphere from the curiosity rover. *Science*, 341(6143):263–266, 2013.

[92] Hugh H. Kieffer. Cold jets in the martian polar caps. *Journal of Geophysical Research: Planets*, 112(E8), 2007.

[93] James W. Head, Harald Hiesinger, Mikhail A. Ivanov, Mikhail A. Kreslavsky, Stephen Pratt, and Bradley J. Thomson. Possible ancient oceans on Mars: evidence from Mars Orbiter Laser Altimeter data. *Science*, 286(5447):2134–2137, 1999.

[94] Michael H. Carr and James W. Head. Geologic history of Mars. *Earth and Planetary Science Letters*, 294(3-4):185–203, June 2010.

[95] Jeffrey C. Andrews-Hanna, Maria T. Zuber, and W. Bruce Banerdt. The borealis basin and the origin of the martian crustal dichotomy. *Nature*, 453(7199):1212–1215, 2008.

[96] Shijie Zhong and Maria T. Zuber. Degree-1 mantle convection and the crustal dichotomy on Mars. *Earth and Planetary Science Letters*, 189(1-2):75–84, June 2001.

[97] Roger J. Phillips, Maria T. Zuber, Sean C. Solomon, Matthew P. Golombek, Bruce M. Jakosky, W. Bruce Banerdt, David E. Smith, Raymond M. Williams, Brian M. Hynek, Oded Aharonson, and II Hauck, Steven A. Ancient geodynamics and global-scale hydrology on Mars. *Science*, 291(5513):2587–2591, 2001.

[98] Michael H. Carr. *The Surface of Mars*. Cambridge University Press, 2006.

[99] Simon C. Stähler, Amir Khan, W. Bruce Banerdt, Philippe Lognonné, Domenico Giardini, Savas Ceylan, Mélanie Drilleau, A. Cecilia Duran, Raphaël F. Garcia, Quancheng Huang, Doyeon Kim, Vedran Lekic, Henri Samuel, Martin Schimmel, Nicholas Schmerr, David Sollberger, Éléonore Stutzmann, Zongbo Xu, Daniele Antonangeli, Constantinos Charalambous, Paul M. Davis, Jessica C. E. Irving, Taichi Kawamura, Martin Knapmeyer, Ross Maguire, Angela G. Marusiak, Mark P. Panning, Clément Perrin, Ana-Catalina Plesa, Attilio Rivoldini, Cédric Schmelzbach, Géraldine Zenhäusern, Éric Beucler, John Clinton, Nikolaj Dahmen, Martin van Driel, Tamara Gudkova, Anna Horleston, W. Thomas Pike, Matthieu Plasman, and Suzanne E. Smrekar. Seismic detection of the martian core. *Science*, 373(6553):443–448, 2021.

[100] Mark A. Wieczorek and Maria T. Zuber. Thickness of the martian crust: Improved constraints from geoid-to-topography ratios. *Journal of Geophysical Research: Planets*, 109(E1), 2004.

[101] Sander Goossens, Terence J. Sabaka, Antonio Genova, Erwan Mazarico, Joseph B. Nicholas, and Gregory A. Neumann. Evidence for a low bulk crustal density for mars from gravity and topography. *Geophysical Research Letters*, 44(15):7686–7694, 2017.

[102] D. Breuer and T. Spohn. Early plate tectonics versus single-plate tectonics on mars: Evidence from magnetic field history and crust evolution. *Journal of Geophysical Research: Planets*, 108(E7), 2003.

[103] M. H. Acuña, J. E. P. Connerney, N. F. Ness, R. P. Lin, D. Mitchell, C. W. Carlson, J. McFadden, K. A. Anderson, H. Rème, C. Mazelle, D. Vignes, P. Wasilewski, and P. Cloutier. Global distribution of crustal magnetization discovered by the Mars Global Surveyor MAG/ER experiment. *Science*, 284(5415):790–793, 1999.

[104] R. J. Lillis, H. V. Frey, and M. Manga. Rapid decrease in martian crustal magnetization in the noachian era: Implications for the dynamo and climate of early mars. *Geophysical Research Letters*, 35(14), 2008.

[105] B. M. Jakosky, J. M. Grebowsky, J. G. Luhmann, J. Connerney, F. Eparvier, R. Ergun, J. Halekas, D. Larson, P. Mahaffy, J. McFadden, D. L. Mitchell, N. Schneider, R. Zurek, S. Bougher, D. Brain, Y. J. Ma, C. Mazelle, L. Andersson, D. Andrews, D. Baird, D. Baker, J. M. Bell, M. Benna, M. Chaffin, P. Chamberlin, Y.-Y. Chaufray, J. Clarke, G. Collinson, M. Combi, F. Crary, T. Cravens, M. Crismani, S. Curry, D. Curtis, J. Deighan, G. Delory, R. Dewey, G. DiBraccio, C. Dong, Y. Dong, P. Dunn, M. Elrod, S. England, A. Eriksson, J. Espley, S. Evans, X. Fang, M. Fillingim, K. Fortier, C. M. Fowler, J. Fox, H. Gröller, S. Guzewich, T. Hara, Y. Harada, G. Holsclaw, S. K. Jain, R. Jolitz, F. Leblanc, C. O. Lee, Y. Lee, F. Lefevre, R. Lillis, R. Livi, D. Lo, M. Mayyasi, W. McClintock, T. McEnulty, R. Modolo, F. Montmessin, M. Morooka, A. Nagy, K. Olsen, W. Peterson, A. Rahmati, S. Ruhunusiri, C. T. Russell, S. Sakai, J.-A. Sauvaud, K. Seki, M. Steckiewicz, M. Stevens, A. I. F. Stewart, A. Stiepen, S. Stone, V. Tenishev, E. Thiemann, R. Tolson, D. Toublanc, M. Vogt, T. Weber, P. Withers, T. Woods, and R. Yelle. Maven observations of the response of mars to an interplanetary coronal mass ejection. *Science*, 350(6261):aad0210, 2015.

[106] R. A. Jacobson. The Orbits and Masses of the Martian Satellites and the Libration of Phobos. AJ, 139(2):668–679, February 2010.

[107] Pascal Rosenblatt. The origin of the Martian moons revisited. A&A Rev., 19:44, August 2011.

[108] Joseph A. Burns. Dynamical characteristics of phobos and deimos. *Reviews of Geophysics*, 10(2):463–483, 1972.

[109] W. K. Hartmann and G. Neukum. Cratering chronology and the evolution of mars. *Space Science Reviews*, 96:165–194, 2001.

[110] D. M. Burr, A. S. McEwen, and S. E. H. Sakimoto. Recent aqueous floods from the cerberus fossae, mars. *Geophysical Research Letters*, 29(13):13–1–13–4, 2002.

[111] Jeffrey J. Plaut, Giovanni Picardi, Ali Safaeinili, Anton B. Ivanov, Sarah M. Milkovich, Andrea Cicchetti, Wlodek Kofman, Jérémie Mouginot, William M. Farrell, Roger J. Phillips, Stephen M. Clifford, Alessandro Frigeri, Roberto Orosei, Costanzo Federico, Iwan P. Williams, Donald A. Gurnett, Erling Nielsen, Tor Hagfors, Essam Heggy, Ellen R. Stofan, Dirk Plettemeier, Thomas R. Watters, Carlton J. Leuschen, and Peter Edenhofer. Subsurface Radar Sounding of the South Polar Layered Deposits of Mars. *Science*, 316(5821):92, April 2007.

[112] Michael D. Smith. The annual cycle of water vapor on mars as observed by the thermal emission spectrometer. *Journal of Geophysical Research: Planets*, 107(E11):25–1–25–19, 2002.

[113] W. C. Feldman, T. H. Prettyman, S. Maurice, J. J. Plaut, D. L. Bish, D. T. Vaniman, M. T. Mellon, A. E. Metzger, S. W. Squyres, S. Karunatillake, W. V. Boynton, R. C. Elphic, H. O. Funsten, D. J. Lawrence, and R. L. Tokar. Global distribution of near-surface hydrogen on mars. *Journal of Geophysical Research: Planets*, 109(E9), 2004.

[114] C. M. Stuurman, G. R. Osinski, J. W. Holt, J. S. Levy, T. C. Brothers, M. Kerrigan, and B. A. Campbell. Sharad detection and characterization of subsurface water ice deposits in utopia planitia, mars. *Geophysical Research Letters*, 43(18):9484–9491, 2016.

[115] Harold Masursky. An overview of geological results from mariner 9. *Journal of Geophysical Research (1896-1977)*, 78(20):4009–4030, 1973.

[116] Timothy A. Goudge, John F. Mustard, James W. Head, and Caleb I. Fassett. Constraints on the history of open-basin lakes on mars from the composition and timing of volcanic resurfacing. *Journal of Geophysical Research: Planets*, 117(E12), 2012.

[117] Briony H. N. Horgan, Ryan B. Anderson, Gilles Dromart, Elena S. Amador, and Melissa S. Rice. The mineral diversity of jezero crater: Evidence for possible lacustrine carbonates on mars. *Icarus*, 339:113526, 2020.

[118] Sharon A. Wilson and James R. Zimbelman. Latitude-dependent nature and physical characteristics of transverse aeolian ridges on mars. *Journal of Geophysical Research: Planets*, 109(E10), 2004.

[119] J. Carter, F. Poulet, J.-P. Bibring, N. Mangold, and S. Murchie. Hydrous minerals on mars as seen by the crism and omega imaging spectrometers: Updated global view. *Journal of Geophysical Research: Planets*, 118(4):831–858, 2013.

[120] Steven W. Squyres, John P. Grotzinger, Raymond E. Arvidson, III Bell, James F., Wendy Calvin, Philip R. Christensen, Benton C. Clark, Joy A. Crisp, William H. Farrand, Kenneth E. Herkenhoff, Jeffrey R. Johnson, Gösta Klingelhöfer, Andrew H. Knoll, Scott M. McLennan, Jr. McSween, Harry Y., Richard V. Morris, Jr. Rice, James W., Ralf Rieder, and Laurence A. Soderblom. In situ evidence for an ancient aqueous environment at meridiani planum, mars. *Science*, 306(5702):1709–1714, 2004.

[121] J. P. Grotzinger, D. Y. Sumner, L. C. Kah, K. Stack, S. Gupta, L. Edgar, D. Rubin, K. Lewis, J. Schieber, N. Mangold, R. Milliken, P. G. Conrad, D. DesMarais, J. Farmer, K. Siebach, F. Calef, J. Hurowitz, S. M. McLennan, D. Ming, D. Vaniman, J. Crisp, A. Vasavada, K. S. Edgett, M. Malin, D. Blake, R. Gellert, P. Mahaffy, R. C. Wiens, S. Maurice, J. A. Grant, S. Wilson, R. C. Anderson, L. Beegle, R. Arvidson, B. Hallet, R. S. Sletten, M. Rice, J. Bell, J. Griffes, B. Ehlmann, R. B. Anderson, T. F. Bristow, W. E. Dietrich, G. Dromart, J. Eigenbrode, A. Fraeman, C. Hardgrove, K. Herkenhoff, L. Jandura, G. Kocurek, S. Lee, L. A. Leshin, R. Leveille, D. Limonadi, J. Maki, S. McCloskey, M. Meyer, M. Minitti, H. Newsom, D. Oehler, A. Okon, M. Palucis, T. Parker, S. Rowland, M. Schmidt, S. Squyres, A. Steele, E. Stolper, R. Summons, A. Treiman, R. Williams, A. Yingst, MSL Science Team, Osku Kemppinen, Nathan Bridges, Jeffrey R. Johnson, David Cremers, Austin Godber, Meenakshi Wadhwa, Danika Wellington, Ian McEwan, Claire Newman, Mark Richardson, Antoine Charpentier, Laurent Peret, Penelope King, Jennifer Blank, Gerald Weigle, Shuai Li, Kevin Robertson, Vivian Sun, Michael Baker, Christopher Edwards, Kenneth Farley, Hayden

Miller, Megan Newcombe, Cedric Pilorget, Claude Brunet, Victoria Hipkin, Richard Léveillé, Geneviève Marchand, Pablo Sobrón Sánchez, Laurent Favot, George Cody, Lorenzo Flückiger, David Lees, Ara Nefian, Mildred Martin, Marc Gailhanou, Frances Westall, Guy Israël, Christophe Agard, Julien Baroukh, Christophe Donny, Alain Gaboriaud, Philippe Guillemot, Vivian Lafaille, Eric Lorigny, Alexis Paillet, René Pérez, Muriel Saccoccio, Charles Yana, Carlos Armiens-Aparicio, Javier Caride Rodríguez, Isaías Carrasco Blázquez, Felipe Gómez Gómez, Javier Gómez-Elvira, Sebastian Hettrich, Alain Lepinette Malvitte, Mercedes Marín Jiménez, Jesús Martínez-Frías, Javier Martín-Soler, F. Javier Martín-Torres, Antonio Molina Jurado, Luis Mora-Sotomayor, Guillermo Muñoz Caro, Sara Navarro López, Verónica Peinado-González, Jorge Pla-García, José Antonio Rodriguez Manfredi, Julio José Romeral-Planelló, Sara Alejandra Sans Fuentes, Eduardo Sebastian Martinez, Josefina Torres Redondo, Roser Urqui-O'Callaghan, María-Paz Zorzano Mier, Steve Chipera, Jean-Luc Lacour, Patrick Mauchien, Jean-Baptiste Sirven, Heidi Manning, Alberto Fairén, Alexander Hayes, Jonathan Joseph, Robert Sullivan, Peter Thomas, Audrey Dupont, Angela Lundberg, Noureddine Melikechi, Alissa Mezzacappa, Julia DeMarines, David Grinspoon, Günther Reitz, Benito Prats, Evgeny Atlaskin, Maria Genzer, Ari-Matti Harri, Harri Haukka, Henrik Kahanpää, Janne Kauhanen, Mark Paton, Jouni Polkko, Walter Schmidt, Tero Siili, Cécile Fabre, James Wray, Mary Beth Wilhelm, Franck Poitrasson, Kiran Patel, Stephen Gorevan, Stephen Indyk, Gale Paulsen, David Bish, Brigitte Gondet, Yves Langevin, Claude Geffroy, David Baratoux, Gilles Berger, Alain Cros, Claude d'Uston, Olivier Forni, Olivier Gasnault, Jérémie Lasue, Qiu-Mei Lee, Pierre-Yves Meslin, Etienne Pallier, Yann Parot, Patrick Pinet, Susanne Schröder, Mike Toplis, Éric Lewin, Will Brunner, Ezat Heydari, Cherie Achilles, Brad Sutter, Michel Cabane, David Coscia, Cyril Szopa, François Robert, Violaine Sautter, Stéphane Le Mouélic, Marion Nachon, Arnaud Buch, Fabien Stalport, Patrice Coll, Pascaline François, François Raulin, Samuel Teinturier, James Cameron, Sam Clegg, Agnès Cousin, Dorothea DeLapp, Robert Dingler, Ryan Steele Jackson, Stephen Johnstone, Nina Lanza, Cynthia Little, Tony Nelson, Richard B. Williams, Andrea Jones, Laurel Kirkland, Burt Baker, Bruce Cantor, Michael Caplinger, Scott Davis, Brian Duston, Donald Fay, David Harker, Paul Herrera, Elsa Jensen, Megan R. Kennedy, Gillian Krezoski, Daniel Krysak, Leslie Lipkaman, Elaina McCartney, Sean McNair, Brian Nixon, Liliya Posiolova, Michael Ravine, Andrew Salamon, Lee Saper, Kevin Stoiber, Kimberley Supulver, Jason Van Beek, Tessa Van Beek, Robert Zimdar, Katherine Louise French, Karl Iagnemma, Kristen Miller, Fred Goesmann, Walter Goetz, Stubbe Hviid, Micah Johnson, Matthew Lefavor, Eric Lyness, Elly Breves, M. Darby Dyar, Caleb Fassett, Laurence Edwards, Robert Haberle, Tori Hoehler, Jeff Hollingsworth, Melinda Kahre, Leslie Keely, Christopher McKay, Lora Bleacher, William Brinckerhoff, David Choi, Jason P. Dworkin, Melissa Floyd, Caroline Freissinet, James Garvin, Daniel Glavin, Daniel Harpold, David K. Martin, Amy McAdam, Alexander Pavlov, Eric Raaen, Michael D. Smith, Jennifer Stern, Florence Tan, Melissa Trainer, Arik Posner, Mary Voytek, Andrew Aubrey, Alberto Behar, Diana Blaney, David Brinza, Lance Christensen, Lauren DeFlores, Jason Feldman, Sabrina Feldman, Gregory Flesch, Insoo Jun, Didier Keymeulen, Michael Mischna, John Michael Morookian, Betina Pavri, Marcel Schoppers, Aaron Sengstacken, John J. Simmonds, Nicole Spanovich, Manuel de la Torre Juarez, Christopher R. Webster, Albert Yen, Paul Douglas Archer, Francis Cucinotta, John H. Jones, Richard V. Morris, Paul Niles, Elizabeth Rampe, Thomas Nolan, Martin Fisk, Leon Radziemski, Bruce Barraclough, Steve Bender, Daniel Berman, Eldar Noe Dobrea, Robert Tokar, Timothy Cleghorn, Wesley

Huntress, Gérard Manhès, Judy Hudgins, Timothy Olson, Noel Stewart, Philippe Sarrazin, Edward Vicenzi, Mark Bullock, Bent Ehresmann, Victoria Hamilton, Donald Hassler, Joseph Peterson, Scot Rafkin, Cary Zeitlin, Fedor Fedosov, Dmitry Golovin, Natalya Karpushkina, Alexander Kozyrev, Maxim Litvak, Alexey Malakhov, Igor Mitrofanov, Maxim Mokrousov, Sergey Nikiforov, Vasily Prokhorov, Anton Sanin, Vladislav Tretyakov, Alexey Varenikov, Andrey Vostrukhin, Ruslan Kuzmin, Benton Clark, Michael Wolff, Oliver Botta, Darrell Drake, Keri Bean, Mark Lemmon, Susanne P. Schwenzer, Ella Mae Lee, Robert Sucharski, Miguel Ángel de Pablo Hernández, Juan José Blanco Ávalos, Miguel Ramos, Myung-Hee Kim, Charles Malespin, Ianik Plante, Jan-Peter Muller, Rafael Navarro-González, Ryan Ewing, William Boynton, Robert Downs, Mike Fitzgibbon, Karl Harshman, Shaunna Morrison, Onno Kortmann, Amy Williams, Günter Lugmair, Michael A. Wilson, Bruce Jakosky, Tonci Balic-Zunic, Jens Frydenvang, Jaqueline Kløvgaard Jensen, Kjartan Kinch, Asmus Koefoed, Morten Bo Madsen, Susan Louise Svane Stipp, Nick Boyd, John L. Campbell, Glynis Perrett, Irina Pradler, Scott VanBommel, Samantha Jacob, Tobias Owen, Hannu Savijärvi, Eckart Boehm, Stephan Böttcher, Sönke Burmeister, Jingnan Guo, Jan Köhler, César Martín García, Reinhold Mueller-Mellin, Robert Wimmer-Schweingruber, John C. Bridges, Timothy McConnochie, Mehdi Benna, Heather Franz, Hannah Bower, Anna Brunner, Hannah Blau, Thomas Boucher, Marco Carmosino, Sushil Atreya, Harvey Elliott, Douglas Halleaux, Nilton Rennó, Michael Wong, Robert Pepin, Beverley Elliott, John Spray, Lucy Thompson, Suzanne Gordon, Ann Ollila, Joshua Williams, Paulo Vasconcelos, Jennifer Bentz, Kenneth Nealson, Radu Popa, Jeffrey Moersch, Christopher Tate, Mackenzie Day, Raymond Francis, Emily McCullough, Ed Cloutis, Inge Loes ten Kate, Daniel Scholes, Susan Slavney, Thomas Stein, Jennifer Ward, Jeffrey Berger, and John E. Moores. A habitable fluvio-lacustrine environment at yellowknife bay, gale crater, mars. *Science*, 343(6169):1242777, 2014.

[122] Geronimo L. Villanueva, Michael J. Mumma, Robert E. Novak, Hans Ulrich Käufl, Paul Hartogh, Thérèse Encrenaz, Alan Tokunaga, Ali Khayat, and Michael D. Smith. Strong water isotopic anomalies in the martian atmosphere: Probing current and ancient reservoirs. *Science*, 348(6231):218–221, 2015.

[123] Michael Carr and James Head. Oceans on mars: An assessment of the observational evidence and possible fate. *J. Geophys. Res. JOURNAL OF GEOPHYSICAL RESEARCH*, 108, 01 1029.

[124] Tomohiro Usui, Conel M. O'D. Alexander, Jianhua Wang, Justin I. Simon, and John H. Jones. Meteoritic evidence for a previously unrecognized hydrogen reservoir on mars. *Earth and Planetary Science Letters*, 410:140–151, 2015.

[125] James Farquhar and Mark H. Thiemens. Oxygen cycle of the martian atmosphere-regolith system: Δ17o of secondary phases in nakhla and lafayette. *Journal of Geophysical Research: Planets*, 105(E5):11991–11997, 2000.

[126] Bruce M. Jakosky and Roger J. Phillips. Mars' volatile and climate history. *Nature*, 412(6843):237–244, 2001.

[127] R. Jordan, G. Picardi, J. Plaut, K. Wheeler, D. Kirchner, A. Safaeinili, W. Johnson, R. Seu, D. Calabrese, E. Zampolini, A. Cicchetti, R. Huff, D. Gurnett, A. Ivanov, W. Kofman,

R. Orosei, T. Thompson, P. Edenhofer, and O. Bombaci. The Mars express MARSIS sounder instrument. Planet. Space Sci., 57(14-15):1975–1986, December 2009.

[128] Roberto Orosei, Rolando Jordan, D. Morgan, Marco Cartacci, A. Cicchetti, Firat Duru, D Gurnett, Essam Heggy, Donald Kirchner, Raffaella Noschese, Wlodek Kofman, A Masdea, J Plaut, Roberto Seu, Thomas Watters, and G Picardi. Mars advanced radar for subsurface and ionospheric sounding (marsis) after nine years of operation: A summary. *Planetary and Space Science*, 112, 11 2014.

[129] Roberto Orosei, Chenglong Ding, Wenzhe Fa, Antonios Giannopoulos, Alain Hérique, Wlodek Kofman, Sebastian E. Lauro, Chao Li, Elena Pettinelli, Yan Su, et al. The global search for liquid water on mars from orbit: Current and future perspectives. *Life*, 10(8):120, 2020.

[130] E. Nielsen, X.-D. Wang, D. A. Gurnett, D. L. Kirchner, R. Huff, R. Orosei, A. Safaeinili, J. J. Plaut, and G. Picardi. Vertical sheets of dense plasma in the topside martian ionosphere. *Journal of Geophysical Research: Planets*, 112(E2), 2007.

[131] Paul Withers. Attenuation of radio signals by the ionosphere of mars: Theoretical development and application to marsis observations. *Radio Science*, 46(2), 2011.

[132] Roberto Seu, Roger J. Phillips, Daniela Biccari, Roberto Orosei, Arturo Masdea, Giovanni Picardi, Ali Safaeinili, Bruce A. Campbell, Jeffrey J. Plaut, Lucia Marinangeli, Suzanne E. Smrekar, and Daniel C. Nunes. Sharad sounding radar on the mars reconnaissance orbiter. *Journal of Geophysical Research: Planets*, 112(E5), 2007.

[133] Donald A. Gurnett, David L. Kirchner, Richard L. Huff, David D. Morgan, Ann M. Persoon, Thomas F. Averkamp, Firdevs Duru, Ewald Nielsen, Ali Safaeinili, Jeffrey J. Plaut, and Giovanni Picardi. Radar soundings of the ionosphere of mars. *Science*, 310(5756):1929–1933, 2005.

[134] Ali Safaeinili, Wlodek Kofman, Jérémie Mouginot, Yonggyu Gim, Alain Herique, Anton B. Ivanov, Jeffrey J. Plaut, and Giovanni Picardi. Estimation of the total electron content of the martian ionosphere using radar sounder surface echoes. *Geophysical Research Letters*, 34(23), 2007.

[135] Andrea Cicchetti, Marco Cartacci, Raffaella Noschese, and Carlo Nenna. Mars express marsis phobos commanding architectures. Technical report, Agenzia Spaziale Italiana, 2016. Accepted in Open Access 2023.

[136] Jacques Laskar, Brigitte Levrard, and John F. Mustard. Orbital forcing of the martian polar layered deposits. *Nature*, 419(6905):375–377, 2002.

[137] Maria T. Zuber, Frank G. Lemoine, David E. Smith, Alex S. Konopliv, Suzanne E. Smrekar, and Sami W. Asmar. Mars reconnaissance orbiter radio science gravity investigation. *Journal of Geophysical Research: Planets*, 112(E5), 2007.

[138] Junlun Li, Jeffrey C. Andrews-Hanna, Youshun Sun, Roger J. Phillips, Jeffrey J. Plaut, and Maria T. Zuber. Density variations within the south polar layered deposits of mars. *Journal of Geophysical Research: Planets*, 117(E4), 2012.

[139] Cyril Grima, Wlodek Kofman, Jérémie Mouginot, Roger J. Phillips, Alain Hérique, Daniela Biccari, Roberto Seu, and Marco Cutigni. North polar deposits of mars: Extreme purity of the water ice. *Geophysical Research Letters*, 36(3), 2009.

[140] C. Mätzler. Microwave Properties of Ice and Snow. In B. Schmitt, C. de Bergh, and M. Festou, editors, *Solar System Ices*, volume 227 of *Astrophysics and Space Science Library*, page 241, January 1998.

[141] G. A. Neumann, J. B. Abshire, O. Aharonson, J. B. Garvin, X. Sun, and M. T. Zuber. Mars orbiter laser altimeter pulse width measurements and footprint-scale roughness. *Geophysical Research Letters*, 30(11), 2003.

[142] M. L. Litvak, I. G. Mitrofanov, A. S. Kozyrev, A. B. Sanin, V. I. Tretyakov, W. V. Boynton, N. J. Kelly, D. Hamara, and R. S. Saunders. Long-term observations of southern winters on mars: Estimations of column thickness, mass, and volume density of the seasonal co2 deposit from hend/odyssey data. *Journal of Geophysical Research: Planets*, 112(E3), 2007.

[143] R. Orosei, S. E. Lauro, E. Pettinelli, A. Cicchetti, M. Coradini, B. Cosciotti, F. Di Paolo, E. Flamini, E. Mattei, M. Pajola, F. Soldovieri, M. Cartacci, F. Cassenti, A. Frigeri, S. Giuppi, R. Martufi, A. Masdea, G. Mitri, C. Nenna, R. Noschese, M. Restano, and R. Seu. Radar evidence of subglacial liquid water on mars. *Science*, 361(6401):490–493, 2018.

[144] Jennifer L. Whitten and Bruce A. Campbell. Lateral continuity of layering in the mars south polar layered deposits from sharad sounding data. *Journal of Geophysical Research: Planets*, 123(6):1541–1554, 2018.

[145] Roger J. Phillips, Brian J. Davis, Kenneth L. Tanaka, Shane Byrne, Michael T. Mellon, Nathaniel E. Putzig, Robert M. Haberle, Melinda A. Kahre, Bruce A. Campbell, Lynn M. Carter, Isaac B. Smith, John W. Holt, Suzanne E. Smrekar, Daniel C. Nunes, Jeffrey J. Plaut, Anthony F. Egan, Timothy N. Titus, and Roberto Seu. Massive co₂ ice deposits sequestered in the south polar layered deposits of mars. *Science*, 332(6031):838–841, 2011.

[146] John D. Toner, David C. Catling, and Bonnie Light. Soluble salts at the phoenix lander site, mars: A reanalysis of the wet chemistry laboratory data. *Geochimica et Cosmochimica Acta*, 136:142–168, 2014.

[147] M. Cartacci, E. Amata, A. Cicchetti, R. Noschese, S. Giuppi, B. Langlais, A. Frigeri, R. Orosei, and G. Picardi. Mars ionosphere total electron content analysis from marsis subsurface data. *Icarus*, 223(1):423–437, 2013.

[148] M. Cartacci, B. Sánchez-Cano, A. Cicchetti, R. Noschese, B. Langlais, and R. Orosei. Mars ionosphere tec estimation from marsis data: A new approach. *Icarus*, 2025.

[149] Giovanni Picardi, S. Sorge, Roberto Seu, J. J. Plaut, W. T. Johnson, Rolando L. Jordan, Donald A. Gurnett, F. Provvedi, E. Zampolini, and Carlo Zelli. MARSIS experiment: design and operations overview. In David A. Noon, Glen F. Stickley, and Dennis Longstaff, editors, *Eighth International Conference on Ground Penetrating Radar*, volume 4084 of *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, pages 228–233, April 2000.

[150] M. Cartacci, B. Sánchez-Cano, R. Orosei, R. Noschese, A. Cicchetti, O. Witasse, F. Cantini, and A. P. Rossi. Improved estimation of mars ionosphere total electron content. *Icarus*, 2017.

[151] Alessandro Piombini. Analisi e mitigazione delle interferenze nei radargrammi marsis: un approccio basato su tecniche di computer vision. Laurea magistrale, Università di Bologna, 2025. Corso di Studio in Ingegneria Informatica [LM-DM270].

[152] Nobuyuki Otsu. A threshold selection method from gray-level histograms. *IEEE Transactions on Systems, Man, and Cybernetics*, 9(1):62–66, 1979.

[153] C. Noviello, G. Esposito, G. Fasano, A. Renga, F. Soldovieri, and I. Catapano. Small-uav radar imaging system performance with gps and cdgps based motion compensation. *Remote Sensing*, 12(20):3463, 2020.

[154] I. Catapano, G. Gennarelli, G. Ludeno, C. Noviello, G. Esposito, and F. Soldovieri. Contactless ground penetrating radar imaging: State of the art, challenges, and microwave tomography-based data processing. *IEEE Geoscience and Remote Sensing Magazine*, 10(1):251–273, March 2022.

[155] G. Pupillo, S. Righini, R. Orosei, C. Bortolotti, G. Maccaferri, M. Roma, M. Mastrogiuseppe, T. Pisanu, L. Schirru, S. Cicalò, A. Tripodo, J. Harju, A. Penttilä, A. K. Virkki, U. Bach, A. Kraus, A. Margheri, R. Ghiani, M. N. Iacolina, and G. Sessler. Toward a european facility for ground-based radar observations of near-earth objects. *Remote Sensing*, 16(1):38, 2024.

[156] NASA Jet Propulsion Laboratory. Near-Earth Object (NEO) Statistics. `https://cneos.jpl.nasa.gov/stats/size.html`, 2026. Accessed: 2026-02-16.

[157] E. F. Helin and E. M. Shoemaker. The Palomar planet-crossing asteroid survey, 1973-1978. Icarus, 40(3):321–328, December 1979.

[158] D. L. Rabinowitz. The Size Distribution of the Earth-approaching Asteroids. ApJ, 407:412, April 1993.

[159] Task Force on Potentially Hazardous Near Earth Objects. Report of the Task Force on Potentially Hazardous Near Earth Objects. Technical report, UK Government, September 2000. Chaired by Sir Crispin Tickell.

[160] Patrick Michel, Vincenzo Zappalà, Alberto Cellino, and Paolo Tanga. Estimated abundance of atens and asteroids evolving on orbits between earth and sun. *Icarus*, 124:371–378, 1996.

[161] William F. Bottke, Alberto Cellino, Paolo Paolicchi, and Richard P. Binzel, editors. *Asteroids III*. University of Arizona Space Science Series. University of Arizona Press, Tucson, dec 2002.

[162] European Space Agency (ESA). Search for comets. `https://neo.ssa.esa.int/search-for-comets`, 2026. ESA Near-Earth Object Coordination Centre (NEOCC). Accessed: 2026-03-10.

[163] Andrea Milani, Giovanni F. Gronchi, Zoran Knežević, M. E. Sansaturio, O. Arratia, L. Denneau, T. Grav, J. Heasley, R. Jedicke, and J. Kubica. Unbiased orbit determination for the next generation asteroid/comet surveys. *Icarus*, 217:355–366, 2012.

[164] Joseph S. Stuart and Richard P. Binzel. Bias-corrected population, size distribution, and impact hazard for the near-earth objects. *Icarus*, 170(2):295–311, 2004.

[165] Alan W. Harris and Germano D'Abramo. The population of near-Earth asteroids. Icarus, 257:302–312, September 2015.

[166] Schelte J. Bus and Richard P. Binzel. Phase II of the Small Main-Belt Asteroid Spectroscopic Survey. The Observations. Icarus, 158(1):106–145, July 2002.

[167] Francesca E. DeMeo, Richard P. Binzel, Stephen M. Slivan, and Schelte J. Bus. An extension of the Bus asteroid taxonomy into the near-infrared. Icarus, 202(1):160–180, July 2009.

[168] Anne Virkki and K. Muinonen. Radar scattering by planetary surfaces modeled with laboratory-characterized particles. *Icarus*, 269, 01 2016.

[169] J.-L. Margot, P. Pravec, P. Taylor, B. Carry, and S. Jacobson. Asteroid systems: Binaries, triples, and pairs. In Patrick Michel, Francesca E. DeMeo, and William F. Bottke, editors, *Asteroids IV*, pages 355–374. University of Arizona Press, Tucson, 2015.

[170] P. Vernazza, R. P. Binzel, C. A. Thomas, F. E. DeMeo, S. J. Bus, A. S. Rivkin, and A. T. Tokunaga. Compositional differences between meteorites and near-earth asteroids. *Nature*, 454(7206):858–860, 2008.

[171] E. Asphaug, C. B. Agnor, and Q. Williams. Hit-and-run planetary collisions. *Nature*, 439(7073):155–160, 2006.

[172] Peter Schulte, Laia Alegret, Ignacio Arenillas, José A. Arz, Penny J. Barton, Paul R. Bown, Timothy J. Bralower, Gail L. Christeson, Philippe Claeys, Charles S. Cockell, Gareth S. Collins, Alexander Deutsch, Tamara J. Goldin, Kazuhisa Goto, José M. Grajales-Nishimura, Richard A. F. Grieve, Sean P. S. Gulick, Kirk R. Johnson, Wolfgang Kiessling, Christian Koeberl, David A. Kring, Kenneth G. MacLeod, Takafumi Matsui, Jay Melosh, Alessandro Montanari, Joanna V. Morgan, Clive R. Neal, Douglas J. Nichols, Richard D. Norris, Elisabetta Pierazzo, Greg Ravizza, Mario Rebolledo-Vieyra, Wolf Uwe Reimold, Eric Robin, Tobias Salge, Robert P. Speijer, Arthur R. Sweet, Jaime Urrutia-Fucugauchi, Vivi Vajda, Michael T. Whalen, and Pi S. Willumsen. The chicxulub asteroid impact and mass extinction at the cretaceous-paleogene boundary. *Science*, 327(5970):1214–1218, 2010.

[173] Clark R. Chapman and David Morrison. Impacts on the earth by asteroids and comets: Assessing the hazard. *Nature*, 367:33–40, 1994.

[174] Olga P. Popova, Peter Jenniskens, Vyacheslav Emel'yanenko, Anna Kartashova, Evgeny Biryukov, Sergey Khaibrakhmanov, Valery Shuvalov, Yuri Rybnov, Alexander Dudorov, Vladimir I. Grokhovsky, Dmitry D. Badyukov, Qing-Zhu Yin, Peter S. Gural, Jim Albers, Mikael Granvik, Lüdo G. Evers, Jan Kuiper, Victor Kharlamov, Andrei Solovyov, Yuri S. Rusakov, et al. Chelyabinsk airburst, damage assessment, meteorite recovery, and characterization. *Science*, 342(6162):1069–1073, 2013.

[175] Steven J. Ostro. Planetary radar astronomy. *Rev. Mod. Phys.*, 65:1235–1279, Oct 1993.

[176] G. H. Pettengill, I. I. Shapiro, M. E. Ash, R. P. Ingalls, L. P. Rainville, W. B. Smith, and M. L. Stone. Radar observations of Icarus. Icarus, 10(3):432–435, May 1969.

[177] R. Scott Hudson and Steven J. Ostro. Shape and non-principal axis spin state of asteroid 4179 toutatis. *Science*, 270(5233):84–86, 1995.

[178] R. S. Hudson, S. J. Ostro, and D. J. Scheeres. High-resolution model of asteroid 4179 toutatis. *Planetary and Space Science*, 51(2):89–101, 2003.

[179] Flaviane C. F. Venditti, Sean E. Marshall, Maxime Devogele, Luisa F. Zambrano-Marin, and Anna McGilvray. The arecibo observatory's legacy and future radar capabilities. *Acta Astronautica*, 210:610–615, September 2023.

[180] Martin A. Slade, Lance A. M. Benner, and Arnold Silva. Goldstone solar system radar observatory: Earth-based planetary mission support and unique science results. *Proceedings of the IEEE*, 99(5):757–769, 2011.

[181] NASA Jet Propulsion Laboratory. DSN telecommunications link design handbook, 810-005, module 101: 70-m subnet telecommunications interfaces. Technical report, Jet Propulsion Laboratory, California Institute of Technology, Pasadena, CA, 2013.

[182] M. Di Martino, S. Montebugnoli, G. Cevolani, S. Ostro, A. Zaitsev, S. Righini, L. Saba, S. Poppi, M. Delbò, A. Orlati, G. Maccaferri, C. Bortolotti, A. Gavrik, and Y. Gavrik. Results of the first Italian planetary radar experiment. Planet. Space Sci., 52(4):325–330, March 2004.

[183] I. Molotov, A. Konovalenko, V. Agapov, A. Sochilina, B. Lipatov, Yu. Gorshenkov, E. Molotov, G. Tuccari, S. Buttaccio, X. Liu, J. Zhang, X. Hong, X. Huang, A. Kus, K. Borkowski, Z. Sika, V. Abrosimov, A. Tsyukh, V. Samodurov, I. Falkovich, and P. Sukhov. Radar interferometer measurements of space debris using the evpatoria rt-70 transmitter. *Advances in Space Research*, 33(3):424–429, 2004.

[184] Defense Express. Ukraine strikes the RT-70 radio telescope in Crimea: Strategic significance and damage assessment, June 24 2024. Accessed: February 25, 2026.

[185] Oliver James White, Guifré Molera Calvés, Shinji Horiuchi, Phil Edwards, Ed Kruzins, Jon Giorgini, Nick Stacy, Andrew Cole, Chris Phillips, and Edwin Peters. Development of radar and optical tracking of near-earth asteroids at the university of tasmania. *Remote Sensing*, 17(3):352, 2025.

[186] S. P. Naidu, L. A. M. Benner, J. L. Margot, M. W. Busch, and P. A. Taylor. Capabilities of earth-based radar facilities for near-earth asteroid observations. *The Astronomical Journal*, 152(4):99, 2016.

[187] Anish Roshi, Nestor Apnote, Esteban Araya, Hector Arce, L. A. Baker, Willem Baan, Tracy Becker, James Breakall, Robert G. Brown, Christiano Brum, Michael Busch, Donald Campbell, Tyler Cohen, Francisco Cordova, Julia Deneva, Maxime Devogele, Timothy Dolch, F. Fernandez-Rodriquez, Tapashree Ghosh, Paul Goldsmith, Leonid Gurvits, Martha Haynes, Carl Heiles, Dylan Hickson, Bret Isham, Robert Kerr, John Kelly, John Kiriazes, Sid Kumar, Jens Lautenbach, M. E. Lebron, N. Lewandowska, Loris Magnani, PK Manoharan, Sean Marshall, Anna McGilvray, A. Mendez, Robert Minchin, Victor Negron, Michael Nolan, Luca Olmi, Flora Paganelli, Nipuni Palliyaguru, Carmen Pantoja, Zsolt Paragi, Stephen

Parshley, Joshua Peek, Benetge Perera, Philip Perillat, Noemi Pinilla-Alonso, Luis Quintero, H. Radovan, Shikha Raizada, Timothy Robishaw, M. Route, Christopher Salter, Alfredo Santoni, Pedrina Santos, Sukanta Sau, Sravani Vaddi, Fabio Vargas, Flaviane Venditti, Arun Venkataraman, Anne Virkki, Amit Vishwas, Sander Weinreb, Dan Werthimer, Alex Wolszczan, and L. F. Zambrano-Maria. The Next Generation Arecibo Telescope (NGAT). In *Bulletin of the American Astronomical Society*, volume 55, page 344, July 2023.

[188] T. J. W. Lazio, K. de Kleer, V. Ravi, et al. The next-generation ground-based planetary radar. Final report, W. M. Keck Institute for Space Studies (KISS), California Institute of Technology, September 2025.

[189] L. A. M. Benner, M. W. Busch, J. D. Giorgini, P. A. Taylor, and J.-L. Margot. Radar observations of near-earth and main-belt asteroids. In Patrick Michel, Francesca E. DeMeo, and William F. Bottke, editors, *Asteroids IV*, pages 165–182. University of Arizona Press, Tucson, 2015.

[190] Stefano Cicalò, Antonio Tripodo, Giuseppe Pupillo, Roberto Orosei, Simona Righini, Tonino Pisanu, Luca Schirru, Jorma Harju, Antti Penttilä, and Anne Virkki. Esa p3-neo-xxii neo observation concepts for radar systems ccn – final study report. Technical Report SSA-P3-NEO-XXII-SDS-CCN-Report, European Space Agency (ESA), 2022. Restricted access.

[191] Shadi Shambayati. Atmosphere attenuation and noise temperature at microwave frequencies. In S. R. Macgregor, editor, *Low-Noise Systems in the Deep Space Network*, chapter 6, pages 255–281. John Wiley and Sons, Hoboken, NJ, USA, 2008.

[192] A. V. Kantak and S. D. Slobin. Atmosphere attenuation and noise temperature models at dsn antenna locations for 1–45 ghz. Technical Report JPL Technical Report 09-14, Jet Propulsion Laboratory, Pasadena, CA, USA, 2009.

[193] Steven J. Ostro. Radar observations of asteroids and comets. *Publications of the Astronomical Society of the Pacific*, 97:877–884, 1985.

[194] European Space Agency. Neo coordination centre (neocc) database. `https://neo.ssa.esa.int/`, 2023. Accessed: 10 October 2023.

[195] E. White, F. D. Ghigo, R. M. Prestage, D. T. Frayer, R. J. Maddalena, P. T. Wallace, J. J. Brandt, D. Egan, J. D. Nelson, and J. Ray. Green Bank Telescope: Overview and analysis of metrology systems and pointing performance. A&A, 659:A113, March 2022.

[196] G. Valente, M. N. Iacolina, R. Ghiani, A. Saba, G. Serra, E. Urru, G. Montisci, S. Mulas, S. W. Asmar, T. T. Pham, et al. The sardinia space communication asset: Performance of the sardinia deep space antenna x-band downlink capability. *IEEE Access*, 10:64525–64534, 2022.

[197] Giuseppe Tuccari, Vladimirs Bezrukovs, and Marina Nechaeva. Digital base band converter as radar vlbi backend. *Latvian Journal of Physics and Technical Sciences*, 49(4):18–29, 2012.

[198] Alan Whitney, Mark Kettenis, Chris Phillips, and Mamoru Sekido. Vlbi data interchange format (vdif). In *Proceedings of the IVS 2010 General Meeting*, Hobart, Australia, February 2010. Accessed: 1 December 2023.

[199] A. U. Tomatic. Iau minor planet electronic circular no. 2021-b127. https://www.minorplanetcenter.net/mpec/K21/K21BC7.html, 2021. 24 January 2021, Minor Planet Center. Accessed: 10 October 2023.

[200] G.B. Marsden. Iau circular no. 3675. IAU Circular, March 1982. Available online: http://www.cbat.eps.harvard.edu/iauc/03600/03675.html (accessed on 10 October 2023).

[201] Marina Brozovic, Steven J. Ostro, Lance A. M. Benner, Jon D. Giorgini, Raymond F. Jurgens, Randy Rose, Michael C. Nolan, Alice A. Hine, Christopher Magri, Daniel J. Scheeres, and Jean-Luc Margot. Radar observations and a physical model of Asteroid 4660 Nereus, a prime space mission target. Icarus, 201(1):153–166, May 2009.

[202] R. F. Jurgens and R. M. Goldstein. Radar Observations at 3. 5 and 12. 6cm Wavelength of Asteroid 433 Eros. Icarus, 28(1):1–15, May 1976.

[203] T. B. Spahr. Iau minor planet electronic circular no. 2005-l19. IAU Minor Planet Electronic Circular, June 2005. Available online: https://minorplanetcenter.net/mpec/K05/K05L19.html (accessed on 10 October 2023).

[204] G. Molera Calvès. *Radio Spectroscopy and Space Science with VLBI Radio Telescopes for Solar System Research*. PhD thesis, Aalto University, Espoo, Finland, 2012. Available online: http://lib.tkk.fi/Diss/2012/isbn9789526045818 (accessed on 27 April 2012).

[205] G. Molera Calvès, S.V. Pogrebenko, G. Cimò, D.A. Duev, T.M. Bocanegra-Bahamòn, J.F. Wagner, J. Kallunki, P. de Vincente, G. Kronschnabl, R. Haas, et al. Observations and analysis of phase scintillation of spacecraft signal on the interplanetary plasma. *Astronomy & Astrophysics*, 564:1–7, 2014.

[206] D. W. E. Green. Iau circular no. 5198. IAU Circular, December 2022. Available online: http://www.cbat.eps.harvard.edu/iau/cbet/005100/CBET005198.txt (accessed on 19 October 2023).

[207] N. R. Draper and H. Smith. *Applied Regression Analysis*. John Wiley & Sons, New York, 2nd edition, 1981.

[208] Michael K. Shepard, Beth Ellen Clark, Michael C. Nolan, Ellen S. Howell, Christopher Magri, Jon D. Giorgini, Lance A. M. Benner, Steven J. Ostro, Alan W. Harris, Brian Warner, Donald Pray, Petr Pravec, Michael Fauerbach, Thomas Bennett, Alain Klotz, Raoul Behrend, Horacio Correia, Josep Coloma, Silvano Casulli, and Andrew Rivkin. A radar survey of M- and X-class asteroids. Icarus, 195(1):184–205, May 2008.

[209] Steven J. Ostro, Robert Connelly, and Leila Belkora. Asteroid shapes from radar echo spectra: A new theoretical approach. *Icarus*, 73(1):15–24, 1988.

[210] Luis A. Santalo. Integral geometry and geometric probability. In Gian-Carlo Rota, editor, *Encyclopedia of Mathematics and Its Applications*, volume 1, pages 1–4. Addison-Wesley, Reading, MA, 1976.

[211] Istvan Maros and Csaba Meszaros. A repository of convex quadratic programming problems. Technical Report DOC 97/6, Department of Computing, Imperial College, London, U.K., 1997.

[212] David L. Mitchell, R. Scott Hudson, Steven J. Ostro, and Keith D. Rosema. Shape of asteroid 433 eros from inversion of goldstone radar doppler spectra. *Icarus*, 124(1):113–133, 1996.

[213] Kenneth Levenberg. A method for the solution of certain non-linear problems in least squares. *Quarterly of Applied Mathematics*, 2(2):164–168, 1944.

[214] Bengt Fornberg. Generation of finite difference formulas on arbitrarily spaced grids. *Mathematics of Computation*, 51(184):699–706, 1988.

[215] Dong C. Liu and Jorge Nocedal. On the limited memory BFGS method for large scale optimization. *Mathematical Programming*, 45(1–3):503–528, 1989.

[216] Gideon E. Schwarz. Estimating the dimension of a model. *The Annals of Statistics*, 6(2):461–464, 1978.

[217] Hirofumi Fujiwara, Toshimitsu Asakura, and Kazumi Murata. On the van Cittert-Zernike theorem. *Optical and Quantum Electronics*, 4(3):197–205, August 1972.

[218] A. Richard Thompson, James M. Moran, and George W. Swenson. *Interferometry and Synthesis in Radio Astronomy.* Springer, Cham, 3 edition, 2017.

[219] J. S. Ulvestad. Space Very Long Baseline Interferometry. In G. B. Taylor, C. L. Carilli, and R. A. Perley, editors, *Synthesis Imaging in Radio Astronomy II*, volume 180 of *Astronomical Society of the Pacific Conference Series*, page 513, January 1999.

[220] P.J. Napier, D.S. Bagri, B.G. Clark, A.E.E. Rogers, J.D. Romney, A.R. Thompson, and R.C. Walker. The very long baseline array. *Proceedings of the IEEE*, 82(5):658–672, 1994.

[221] Jacob Kooi, Melissa Soriano, James Bowen, Zubair Abdulla, Lorene Samoska, Andy Fung, Raju Manthena, Daniel Hoppe, Hamid Javadi, Timothy Crawford, Darren Hayton, Malo-Gómez Inmaculada, Juan Daniel Gallego Puyol, Ahmed Akgiray, Bekari Gabritchidze, Kieran Cleary, Christopher Jacobs, and Joseph Lazio. A multioctave 8 ghz−40 ghz receiver for radio astronomy. *IEEE Journal of Microwaves*, PP:1–17, 04 2023.

[222] G. B. Taylor, C. L. Carilli, and R. A. Perley, editors. *Synthesis Imaging in Radio Astronomy II*, volume 180 of *ASP Conference Series.* Astronomical Society of the Pacific, 1999.

[223] A. Segalovitz and B. R. Frieden. A 'CLEAN'-type Deconvolution Algorithm. A&A, 70:335, November 1978.

[224] T. J. Pearson and A. C. S. Readhead. Image Formation by Self-Calibration in Radio Astronomy. ARA&A, 22:97–130, January 1984.

[225] Michael W. Busch, Shrinivas R. Kulkarni, Walter Brisken, Steven J. Ostro, Lance A. M. Benner, Jon D. Giorgini, and Michael C. Nolan. Determining asteroid spin states using radar speckles. *Icarus*, 2010.

[226] V. A. Alekseev, B. N. Lipatov, and V. A. Reznikova. Radar-vlbi method: The analysis of its capabilities for determination of planet-rotation parameters. *Radiophysics and Quantum Electronics*, 43:607–613, August 2000.

[227] I. Molotov, V. Agapov, V. Titenko, Z. Khutorovsky, Yu. Burtsev, I. Guseva, V. Rumyantsev, M. Ibrahimov, G. Kornienko, A. Erofeeva, V. Biryukov, V. Vlasjuk, R. Kiladze, R. Zalles, P. Sukhov, R. Inasaridze, G. Abdullaeva, V. Rychalsky, V. Kouprianov, O. Rusakov, and E. Filippov. International scientific optical network for space debris research. *Advances in Space Research*, 40(9):1282–1290, 2007.

[228] M. Nechaeva, A. Antipenko, V. Bezrukovs, D. Bezrukov, A. Dementjev, N. Dugin, A. Konovalenko, V. Kulishenko, X. Liu, A. Nabatov, V. Nesteruk, G. Pupillo, A. Reznichenko, E. Salerno, I. Shmeld, O. Shulga, Y. Sybiryakova, Y. Tikhomirov, A. Tkachenko, A. Volvach, and W.-J. Yang. An experiment on radio location of objects in the near-earth space with vlbi in 2012. *Baltic Astronomy*, 22:341–346, 2013.

[229] M. Nechaeva, A. Antipenko, D. Bezrukov, Vl. Bezrukovs, A. Dementjev, N. Dugin, N. Jekabsons, R. Khutornoy, M. Klapers, A. Konovalenko, V. Kulishenko, A. Nabatov, V. Nesteruk, G. Pupillo, A. Reznichenko, E. Salerno, I. Shmeld, K. Skirmante, Y. Tikhomirov, and V. Voytyuk. First results of the vlbi experiment on radar location of the asteroid 2012 da14. *Baltic Astronomy*, 22, 2013. Received 2013 September 3; accepted 2013 September 17.

[230] European Space Agency (ESA). Close approach of asteroid 2025 fa22. `https://www.esa.int/ESA_Multimedia/Images/2025/09/Close_approach_of_asteroid_2025_FA22`, 2025. Accessed: 2026-03-04; ESA Space Safety article on close approach of asteroid 2025 FA22, published 16 September 2025.

[231] L.A.M. Benner. Goldstone radar observations planning: 2025 fa22, 2025 rl2, 1998 fw4, 2022 sw12, and 2025 pj1. `https://echo.jpl.nasa.gov/asteroids/2025fa22.2025.goldstone.planning.html`, 2025.

[232] European VLBI Network. Evnstatus.txt: Evn status tables. `https://www.evlbi.org/sites/default/files/shared/EVNstatus.txt`, 2024. EVN status table last updated 29 April 2024; provides observatory and telescope availability information for the European VLBI Network.

[233] T. M. Bocanegra-Bahamón, G. Molera Calvés, L. I. Gurvits, D. A. Duev, S. V. Pogrebenko, G. Cimò, D. Dirkx, and P. Rosenblatt. Planetary Radio Interferometry and Doppler Experiment (PRIDE) technique: A test case of the Mars Express Phobos Flyby. II. Doppler tracking: Formulation of observed and computed values, and noise budget. A&A, 609:A59, January 2018.